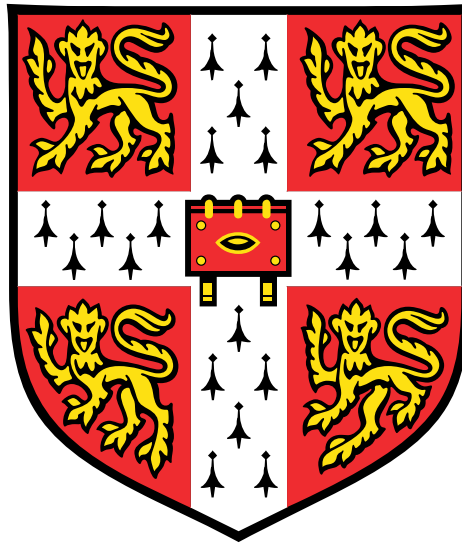# Essays in International Macroeconomics and Finance

Samuel Mann

Faculty of Economics

University of Cambridge

This dissertation is submitted for the degree of

Doctor of Philosophy

Wolfson College   March 2018

To Sofina

Declaration

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration, except as specified in the text. It 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. I further state that no substantial part of my dissertation 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. It does not exceed the prescribed word limit of 60,000 words.

Samuel Mann

March 2018

Acknowledgements

I am deeply grateful to my supervisor, Giancarlo Corsetti, and my adviser, Alexey Onatski, for their continued guidance and support. Working with my co-author, João Duarte, has been a truly enriching experience. Other faculty members at the University of Cambridge have helped me develop my ideas: Tiago Cavalcanti, Chryssi Giannitsarou, Edoardo Gallo, Donald Robertson, Alexander Rodnyansky, Pontus Rendahl, Vasco Carvalho and Petra Geraats. Moreover, I gratefully acknowledge the Faculty's financial support.

At the Cambridge Faculty of Economics, numerous friendships have enriched my Ph.D. life beyond measure. I will never forget the hours spent with Simon Lloyd, David Minarsch, River Chen, Margit Reischer, Dan Wales, Jeroen Dalderop, Felix Grey, Anil Ari, Monica Petrescu, Jasmine Xiao and Farid Ahmed. Dominik Bilgeri, Felix Karstens and Klemens Jansen, as well as many other friends far and wide, have all played their part outside of the faculty environment. Special thanks go to André Meier, who has continually supported me as a friend, mentor and role-model.

This thesis would not exist without the encouragement of my parents, Eva and Christian. They instilled in me a scientific mindset and at the same time taught me to believe in the wonders of our world. I am grateful to my sister, Anna, who has always provided me with helpful advice, as well as her partner, Christian, who supported me during my first ever economic research project at the University of St. Gallen.

Most of all, I am indebted to my fiancée, Sofina, who has been there unconditionally from day one, never failing to encourage, support and motivate me.

Abstract

This collection of essays examines the topic of macroeconomic stabilisation in an international context, focusing on monetary policy, capital controls and exchange rates.

Chapter 1, written in collaboration with Giancarlo Corsetti and João Duarte, reconsiders the effects of common monetary policy shocks across countries in the euro area, using a data-rich factor model and identifying shocks with high-frequency surprises around policy announcements. We show that the degree of heterogeneity in the response to shocks, while being low in financial variables and output, is significant in consumption, consumer prices and macro variables related to the labour and housing markets. Mirroring country-specific institutional and market differences, we find that home ownership rates are significantly correlated with the strength of the housing channel in monetary policy transmission. We document a high dispersion in the response to shocks of house prices and rents and show that, similar to responses in the US, these variables tend to move in different directions.

In Chapter 2, I build a two-country, two-good model to examine the welfare effects of capital controls, finding that under certain circumstances, a shut-down in asset trade can be a Pareto improvement. Further, I examine the robustness of the result to parameter changes, explore a wider set of policy instruments and confront computational issues in this class of international macroeconomic models. I document that within an empirically relevant parameter span for the trade elasticity, the gains from capital controls might be significantly larger than suggested by previous contributions. Moreover, I establish that a refined form of capital controls in the shape of taxes and tariffs cannot improve upon the outcome under financial autarky. Finally, results show that the conjunction of pruning methods and endogenous discount factors can remove explosive behaviour from this class of models and restore equilibrating properties.

In Chapter 3, I use a panel of 20 emerging market currencies to assess whether a model that combines fundamental and non-fundamental exchange rate forecasting approaches can successfully predict risk premia (i.e. currency excess returns) over the short horizon. In doing so, I aim to overcome three main shortcomings of earlier research: i) Sensitivity to the chosen sample period; ii) seemingly arbitrary selection of explanatory variables that differs from currency to currency; and iii) difficulty in interpreting forecasts beyond the numerical signal. Based on a theoretical model of currency risk premia, I use real exchange rate strength combined with indicators for carry, momentum and economic sentiment to homogeneously forecast risk premia across all 20 currencies in the sample at a monthly frequency. In doing so, the model remains largely agnostic about structural choices, keeping arbitrarily imposed restrictions to a minimum. Results from portfolio construction suggest that returns are significant and robust both across currencies as well as over time, with Sharpe Ratios in out-of-sample tests above 0.7.

Contents

3   Bridging Fundamental and Market Sentiment Approaches in Emerging Market Currency Valuation   **83**

Chapter 1

One Money, Many Markets: A Factor Model Approach to Monetary Policy in the Euro Area with High-Frequency Identification

This chapter is co-authored with Giancarlo Corsetti and João Duarte.

1.1 Introduction

Monetary policy in the euro area has long been challenged by financial, economic and institutional heterogeneity among member countries. Although over time there has been some convergence in financial markets, the convergence process has slowed down markedly since the financial crisis (see ECB, 2017). Other markets have remained remarkably different across member countries. Most notably, the institutional backgrounds in labour and housing are highly dissimilar across the currency block.

Two longstanding questions are to which extent these differences in institutional backgrounds imply a heterogeneous transmission of the European Central Bank's (ECB) common monetary policy and, in particular, which specific institutional characteristics drive the observed heterogeneity. These questions are of first-order importance from a policy perspective.[1] Naturally, the ECB would benefit from knowing how national policies and reforms of the institutional framework in a particular economy affect monetary transmission. At the same time, national policy makers would gain from understanding the implications of their policies and reforms for the transmission of monetary policy.

---

[1]The early policy and empirical debate on this issue is summarized by Angeloni et al. (2003), see also Berben et al. (2004).

In this paper, we investigate heterogeneity in the transmission of monetary policy across the euro area (EA) using a dynamic factor model (DFM), and take a first step towards relating the observed heterogeneity to cross-border differences in institutions and markets. We assemble a large dataset including economic and financial time series for the EA as a block and for the 11 original member countries, spanning the years from 1999 to 2016. The high dimensionality of the data allows us to carry out a formal comparison of the degree of heterogeneity among responses to monetary policy shocks across different dimensions of the economy, such as output and asset prices, as well as housing and labour markets. To identify monetary policy shocks, we construct an external instrument using high-frequency changes in asset prices around ECB policy announcements, following the contributions by Gurkaynak et al. (2005) and Gertler and Karadi (2015). Comparing country-specific institutional characteristics in national housing markets as a case study, we show that these characteristics are strongly correlated to the strength of the housing channel in monetary policy transmission.

Our main results are as follows. First, at the aggregate EA level, we find that the factor model results are in line with theory and, notably, that the transmission of monetary shocks does not suffer from the price puzzle. Second, we show that the estimated country-level effects are significantly heterogeneous in prices and variables related to labour and housing markets—some of the least integrated markets in the euro area. The degree of heterogeneity among responses to policy is instead low in financial variables and output. Third, we use our comparative analysis of European housing markets to show that the strength of the housing channel is correlated with differences in home ownership rates, which, as we argue, reflect different institutional characteristics across euro area countries.

On methodological ground, the paper's main contribution lies in the construction of the external instrument series, which we base on changes in the 1-year Euro Overnight Index Average (EONIA) swap rate (i.e. the Overnight Index Swap (OIS) rate for the euro area) around policy announcements. In doing so, we overcome major data availability issues by combining intraday data with end-of-day data from different timezones, creating de-facto intraday series where actual intraday data is unavailable.[2] This solution for the construction of our instrument is not only mechanically feasible, but also economically meaningful, as it highlights the implications of using various means of policy communication—press releases, press statements and Q&A sessions—for the transmission of current and expected future policy. The approach helps us to create a broad measure of monetary policy surprises that incorporates all of the communication channels above. Finally, we test for the relevance of the series in a small VAR, confirming its

---

[2]Intraday data on EONIA swaps is only available for recent years. However, we were able to combine end-of-day data from Tokyo and London to create a de-facto intraday series that goes back to the introduction of the euro. We then compared a narrowly constructed instrument over a sub-sample for which we had complete intraday data with our proposed de-facto intraday series. We find that the series are not significantly different for the sub-sample. See Section 1.2.3 for details.

validity as an external instrument. At the time of writing and to the best of our knowledge, there are very few attempts to construct an external instrument for EA monetary policy.[3] A notable exception is Jarocinski and Karadi (2018). In contemporaneous work, these authors use the high-frequency co-movement of interest rates and stock prices around a narrow window of the policy announcement to disentangle policy from information shocks. Notably, the effects of the monetary shocks we identify in this paper are close to the effects of policy shocks (as opposed to information shocks) documented by these authors.

Our second contribution is motivated by the need to test heterogeneity in the responses of economic variables to a common shock. For each set of impulse responses (e.g. GDP across member countries), we calculate the coefficient of variation statistic, also known as relative standard deviation. The coefficient of variation for a variable is defined as the standard deviation of responses across countries with respect to the EA response, normalised by the size of the EA response. This yields a statistical measure of the dispersion of impulse responses which allows us to carry out comparisons across variables. We employ a bootstrapping procedure to obtain error bands for the coefficient of variation of each variable as well as pairwise differences across variables.

In specifying our empirical model, we build on two strands of the literature. The first is factor modelling, which has emerged as early as the 1970s[4] and was more recently popularised for monetary policy analysis by Bernanke et al. (2005). These authors model macroeconomic interaction with a factor-augmented VAR (FAVAR) that combines factors and perfectly observable series, typically interest rates, in one dynamic system. As a special case of FAVARs, dynamic factor models (DFMs) only contain unobservable factors. Among other applications of DFMs, Stock and Watson (2012) use this approach to disentangle the channels of the 2007-09 recession.

From an applied perspective, the prime advantage of a factor approach[5] is its ability to keep track of individual country-level responses to a common monetary policy shock without heavy parameterisation. Looking at the alternatives, country-by-country VARs incur the cost of heavy parameterisation, while a large panel VAR (PVAR) with all countries imposes restrictions on the individual dynamics. The dynamic factor model solves both problems and provides dynamic effects on the individual countries—including net spillovers—while keeping the parameter space small. In addition, the assumptions on the information structure in the dynamic factor model naturally fit the EA setting. The ECB follows not only a large number of euro-wide series but

---

[3]Limitations on the availability of data are a strict constraint to construct this instrument for the euro area. Some work, such as Kim and Other (2017), resorts to daily data. However, as the resulting windows around policy announcements become very large, they are more likely contaminated by shocks other than those stemming from monetary policy. Notably, European data is frequently released on the morning of a policy meeting, leading to a systematic disturbance of instruments that rely on daily data.

[4]see Stock and Watson (2016) for a comprehensive exposition of factor models, including their early history.

[5]see e.g. Giannone et al. (2005), Bernanke et al. (2005), Stock and Watson (2005) and Forni and Gambetti (2010).

also series in individual member countries. Hence, an empirical model with a small number of variables that does not include country-level data is unlikely to span the information set used by the ECB.[6]

While we closely follow the methodology of Stock and Watson (2012) when we construct our DFM, we bridge the approach with developments from the literature on high-frequency identification and external instruments. As is well known, estimations of monetary policy transmission suffer from an identification problem. One common way to overcome this problem and identify monetary policy shocks is to impose additional internal structure on the VAR, such as timing or sign restrictions. Alternatively, one can add information from outside of the VAR, termed an external instrument approach. We make use of the latter. The two leading examples of existing external instruments for monetary policy shocks in the US are the Romer and Romer (2002) instrument based on a narrative approach, and the high-frequency approach by Gurkaynak et al. (2005). For the second, the key idea is that by choosing a narrow time window around policy announcements, any surprises occurring within the window are most likely only associated with monetary policy shocks. The idea to use high-frequency changes in asset prices, specifically interest rate derivatives, has also been developed by Kuttner (2001), Hamilton (2008) and Campbell et al. (2012).[7] Building on these contributions, Gertler and Karadi (2015) identify monetary policy shocks in a VAR using high frequency changes in Fed funds futures. This paper builds a hybrid between the high-frequency identification proposed by Gertler and Karadi (2015) and the dynamic factor model of Stock and Watson (2012).

The analysis of the housing channel conducted in our paper is related to Calza et al. (2013), who studied how heterogeneity in the structure of housing finance across member countries in the euro area can affect the transmission of monetary policy to house prices. Differently from their work, we take a more comprehensive approach and, specifically, document how differences in home ownership rates are closely linked to asymmetries in house price responses. Moreover, we investigate the role of rents in the housing channel and show that together with house prices, they have a strong link to responses in consumption. More generally, our work is related to the large body of policy and academic research that, given the importance of the topic, has been devoted to the heterogeneous transmission of monetary policy across EA member states. Among the leading examples are Ciccarelli et al. (2013), who look at heterogeneity from the perspective of financial fragility, as well as Barigozzi et al. (2014) who, similar to the methodology followed in this paper, rely on a factor model, although identifying shocks with sign restrictions and pursuing a less comprehensive study, both in the number of variables included and the methodological

---

[6]Other seminal contributions on dynamic factor modelling include Sargent and Sims (1977), Sargent (1989), Giannone et al. (2005) and Boivin and Giannoni (2007).

[7]Further applications of high-frequency identification in the context of monetary policy can be found in Hanson and Stein (2015), Nakamura and Steinsson (2013), Bagliano and Favero (1999), Cochrane and Piazzesi (2002), Faust et al. (2004) and Barakchian and Crowe (2013), among others.

and empirical questions addressed.

In the next section, we describe the methodology used in the empirical analysis and provide details on the external instrument used for the identification of monetary policy shocks. In Section 1.3, we present our results, tracing out the effects of monetary policy on the EA as a whole, as well as on individual member countries. Section 1.4 uses the housing market as a case study to uncover how institutional differences are affecting the monetary transmission across the euro area. Section 1.5 concludes.

1.2   A Dynamic Factor Model for the Euro Area

We begin by motivating the use of a dynamic factor model for the EA and laying out the empirical framework. Later in this section we provide details about the external instrument we construct to identify monetary policy shocks. At the end of the section, we discuss the large data set and estimation.

1.2.1   Motivation

Given the EA setting, we are fundamentally interested in studying the effects of a common monetary policy shock on the EA as a block and on its individual member countries.[8] Recovering both the effects on the block and on member countries imposes some empirical challenges and trade-offs. On the one hand, fully recovering the effects of monetary policy on each individual member country comes with heavy parameterisation. On the other hand, reducing the parameter space by imposing restrictions prevents us from studying the full width of heterogeneous effects. In addition, a small data sample in the time dimension, as encountered in the context of the EA, further increases the acuteness and relevance of this trade-off.

We propose a dynamic factor model for the EA as a parsimonious way to avoid heavy parameterisation while keeping track of individual country responses to the common monetary policy shock. The dynamic factor model allows us to capture dynamic effects on individual countries through unobservable common components. The dimensionality reduction achieved through the factor model allows us to get statistically robust dynamic effects on the individual countries while keeping the parameter space small.

The dynamic factor model has another set of appealing features for the EA. Firstly, we can relax the informational assumption that both the ECB and the econometrician perfectly observe all relevant economic variables. Secondly, as the ECB monitors a large number of indicators in the process of policy formulation, including on country level, it is necessary for the econometrician to take account of the same information set. The DFM achieves this. Finally,

---

[8]A similar setting would appear if one was simultaneously interested in the effects of monetary policy on the U.S. as a whole and at the individual State level.

the dynamic factor model provides a format that is consistent with economic theory. We next address each of these points.

In using a dynamic factor model we do not have to take a stand on specific observable measures corresponding to theoretical concepts. This point was convincingly put forward by Bernanke et al. (2005). In the EA context, this relaxation becomes more relevant as it is harder to find observable Eurowide variables—often weighted averages of individual member countries—that correspond to concepts of economic theory. For example, the concept of *economic activity* in the EA may not be perfectly measured by taking a weighted average of real GDP across countries, given compositional changes that cannot be captured by treating the EA as a single economy in a theoretical model.

The European Central Bank follows not only a large number of eurowide series but also a large number of individual member countries' series. Hence, an empirical model with a small number of variables that does not include country data is unlikely to span the information set used by the ECB. This naturally motivates the inclusion of country-level series in our analysis.

The state-space representation of the dynamic factor model also provides a clear link with economic theory, which creates the opportunity to formally test different mechanisms aimed at explaining the dynamic effects found in this paper. Moreover, given the large size of dynamic effects found in observables, it is possible to test interactions of different mechanisms using the same model and dataset.

There are alternatives to the DFM approach chosen by us—notably Panel VAR and Global VAR models. Both of these approaches involve restricting or explicitly modelling the dynamics through which variables in different units affect each other. These restrictions come at a cost of higher parameterisation relative to the dynamic factor model. Given that we are not explicitly interested in these interaction at the cross-sectional level, but rather in the final net effect, we choose the dynamic factor model for efficiency gains. Ciccarelli et al. (2013) provide a further insightful discussion of the differences between these three approaches.

1.2.2 Empirical Framework

We consequently use the DFM to model macroeconomic interaction. In doing so, we largely follow the methodology proposed by Stock and Watson (2012).

Given a vector of n macroeconomic series $X_t = (X_{1t}, ..., X_{nt})'$ we first model each series as a combination of factors and idiosyncratic disturbances:

$$X_t = \Lambda F_t + e_t, \tag{1.1}$$

where F_t is a vector of unobserved factors, Λ is an $n \times r$ matrix of factor loadings and $e_t = (e_{1t}, ..., e_{nt})'$ denotes a vector of n disturbances. We can interpret ΛF_t as the 'common component' of X_t, whilst e_t is the 'idiosyncratic component'. The evolution of factors is characterised by the following VAR:

$$F_t = \Phi_1 F_{t-1} + \Phi_2 F_{t-2} + ... + \Phi_s F_{t-s} + \eta_t, \tag{1.2}$$

which can be rewritten with lag-operator notation as

$$\Phi(L)F_t = \eta_t, \tag{1.3}$$

where $\Phi(L)$ is a $p \times r$ matrix of lag polynomials and η_t a vector of r innovations. This equation characterises all dynamics in the model. As it stems solely from the interaction of factors, there is no need to model the co-movement of observed variables, hence avoiding the curse of dimensionality.

The static factors can be estimated by suitable cross-sectional averaging. Whilst a setup with multiple factors and general factor loadings does not allow for simple cross-sectional averaging to produce a consistent estimate of the factors, the idea can be generalised using principal components analysis. Given large n and T, the principal components approach estimates the space spanned by the factors, even though the factors themselves are not estimated consistently. Put differently, F_t is estimated consistently up to premultiplication by an arbitrary nonsingular $r \times r$ matrix. The resulting normalisation problem can be resolved by imposing the restriction that $\Lambda'\Lambda = I_r$. Given that this restriction is chosen arbitrarily, the factors cannot be directly interpreted in an economic sense. For most parts, we will work with the reduced-form DFM, making the normalisation inconsequential.

More generally, principal component analysis provides the factors that explain the most variation in the data, while at the same time avoiding an information overlap between the factors as they are orthogonal to each other.

1.2.3 Identification

This section turns to the identification of the monetary policy shocks in the DFM. As is well known, estimations of monetary policy suffer from an identification problem, as monetary policy contemporaneously reacts to other variables in the model. To find the part of the variation in monetary policy that is orthogonal to other variables, various approaches have been proposed in the literature. In traditional VAR-type models, researchers have typically imposed some internal structure on the coefficients in the VAR, such as timing restrictions or sign restrictions.

More recently, Montiel Olea et al. (2012) as well as others have proposed an additional method, where information from outside the VAR is used to identify monetary policy. In the so-called external instrument approach, an instrument is employed that is correlated with the structural shock that the researcher tries to uncover, while being uncorrelated with all other shocks in the system. This corresponds to the standard assumptions of relevance and exogeneity in the instrumental variables literature.

The main concept behind using an external instrument is that when regressing the VAR innovations η_t on the instrument Z_t, the fitted value of the regression identifies the structural shock—up to sign and scale. In fact, as this approach uncovers the covariance between η_t and Z_t, a regression of the instrument on the VAR innovations would equally uncover the structural shock.

Following the VAR literature and the notation in Stock and Watson (2012), we model a linear relationship between the VAR innovations η_t and the structural shocks ϵ_t:

$$\eta_t = H\epsilon_t = [H_1 \cdots H_r] \begin{pmatrix} \epsilon_{1t} \\ \vdots \\ \epsilon_{rt} \end{pmatrix}, \tag{1.4}$$

where H is a matrix of coefficients and H_1 is the first column of H. It follows that $\Sigma_{\eta\eta} = H\Sigma_{\epsilon\epsilon}H'$, with $\Sigma_{\eta\eta} = E(\eta_t\eta_t')$ and $\Sigma_{\epsilon\epsilon} = E(\epsilon_t\epsilon_t')$. If the system is invertible—a standard assumption in the VAR literature—structural shocks can be expressed as linear combinations of innovations:

$$\epsilon_t = H^{-1}\eta_t. \tag{1.5}$$

The main interest in the DFM, as in other VAR-type models, lies in uncovering impulse response functions (IRFs) to a specific shock. To find the impulse response function of X_t with respect to the i^{th} structural shock, we can use equations 1.3 and 1.5 to get

$$F_t = \Phi(L)^{-1}H\epsilon_t. \tag{1.6}$$

Substituting 1.6 into 1.1, we find that

$$X_t = \Lambda\Phi(L)^{-1}H\epsilon_t + e_t. \tag{1.7}$$

where the IRF is $\Lambda\Phi(L)^{-1}H$. Λ and $\Phi(L)$ are already identified from the reduced form, equation 1.2, which we can estimate via ordinary least squares. However, this leaves the identification of H_t, which is dealt with in the next section.

As mentioned above, we identify the shock of interest, say ϵ_{1t}, using the instrumental variable Z_t. The necessary conditions are:

1. Relevance: $E(\epsilon_{1t}Z_t) = \alpha \neq 0$

2. Exogeneity: $E(\epsilon_{jt}Z_t) = 0$, $j = 2, ..., r$

3. Uncorrelated shocks: $\Sigma_{\epsilon\epsilon} = D = diag(\sigma^2_{\epsilon_1}, ..., \sigma^2_{\epsilon_r})$,

where D is an $r \times r$ matrix. The last condition is the standard structural VAR assumption that structural shocks are uncorrelated. This assumption does not fix the variance of shocks. From equation 1.4 we get

$$E(\eta_t Z_t) = E(H\epsilon_t Z_t) = (H_1 \cdots H_r) \begin{pmatrix} E(\epsilon_{1t}Z_t) \\ \vdots \\ E(\epsilon_{rt}Z_t) \end{pmatrix} = H_1 \alpha, \tag{1.8}$$

where the last identity follows from the relevance and exogeneity conditions. It follows that H_1 is identified up to scale and sign by the covariance between the VAR innovations and the instrument. To identify the shocks themselves, we need the third condition on uncorrelated shocks. It implies that we can rewrite the varianance-covariance matrix of η_t as

$$\Sigma_{\eta\eta} = H\Sigma_{\epsilon\epsilon}H' = HDH'. \tag{1.9}$$

Moreover, defining by Π the matrix of coefficients from the population regression of Z_t on η_t, the fitted value of this regression is

$$\Pi\eta_t = E(Z_t\eta_t')\Sigma_{\eta\eta}^{-1}\eta_t, \tag{1.10}$$

which, using equation 1.8 and 1.9, can be written as

$$E(Z_t\eta_t')\Sigma_{\eta\eta}^{-1}\eta_t = \alpha H_1'(HDH')^{-1}\eta_t. \tag{1.11}$$

By simplifying and using equation 1.5, we obtain

$$\alpha H_1'(HDH')^{-1}\eta_t = \alpha(H_1'(H')^{-1})D^{-1}\epsilon_t. \tag{1.12}$$

Finally, we note that $H^{-1}H_1 = e_1$, where $e_1 = (1, 0, ..., 0)'$, which implies that

$$\alpha(H_1'(H')^{-1})D^{-1}\epsilon_t = (\alpha/\sigma^2_{\epsilon_1})\epsilon_{1t} = \Pi\eta_t. \tag{1.13}$$

This conforms with the original statement that the fitted value of a regression of the instrument on the innovations, i.e. $\Pi\eta_t$, identifies the structural shock ϵ_{1t} up to a constant. For additional intuition, Stock and Watson (2012) point out that if the structural shocks ϵ_t were observable and we could hence regress the instrument on the structural shocks, the predicted value would again uncover the shock ϵ_{1t}, up to scale, as the coefficients on all other elements of ϵ_t would be zero. This follows from the relevance and exogeneity conditions of the instrument. Equation 1.13 shows that the projection of Z_t on η_t provides the exact same result, uncovering ϵ_{1t}. Note that to estimate the structural shock, we use the sample analogue of the above equation.

Instrument - "Scripta Volant, Verba Manent"[9]

To obtain an instrument that fulfills the necessary requirement of only being correlated with the monetary policy shock, we build a new series of high frequency surprises around ECB policy announcements. A similar approach has previously been proposed for US monetary policy by Gertler and Karadi (2015). The key idea is that by choosing a narrow time window around policy announcements, any surprises occurring within the window are most likely only associated with monetary policy shocks. Put differently, the assumption is that no other major structural shocks occur during the chosen window around the policy announcement. Correspondingly, all endogenous monetary policy, i.e. all expected monetary policy, is assumed to already have been priced in before the window starts. Consequently, endogenous monetary policy would not cause a change in the instrument at the time of the announcement.

For the instrument we choose changes in the 1-year Euro Overnight Index Average (EONIA) swap rate. The logic goes that while expectations about future policy rate changes are already priced in, unexpected policy shocks will cause the swap to appreciate or depreciate instantly. If market participants, for example, expect a hike in the policy rate by a certain amount, the announcement of such a hike will not cause the 1-year EONIA swap rate to move. However, should a hike or cut be out of line with expectations, the swap rate will adjust as soon as the announcement is made. Similarly, any policy action that changes expectations about future rate movements—often termed 'forward guidance'—will have an impact on the swap. Lloyd (2017a) and Lloyd (2017b) demonstrates that 1 to 24-month Overnight Indexed Swap (OIS) rates accurately measure interest rate expectations. As our chosen EONIA swap rate is the corresponding OIS rate for the euro area, this finding is directly applicable to our instrument, allowing us to capture not only current monetary policy, but also expectations about the future path of monetary policy.

---

[9]The original quotation (*Verba volant, scripta manent*), attributed to Caius Titus, roughly translates as "spoken words fly away, written words remain." We find that, on the contrary, it is often the spoken word of the ECB President during the press conference and Q&A session, which has a larger impact on markets than the written word of the monetary policy press release.

When deciding on the tenor of the EONIA swap, two considerations have to be taken into account. Firstly, to capture how a monetary policy shock affects interest rates across the whole yield curve, a longer dated swap is better suited compared to one with a shorter tenor. On the other hand, however, term premia play a larger role at longer horizons, potentially contaminating the information about future short rates. In dealing with this trade-off, we choose the 1-year rate, based on the observation that 1-year rates are highly sensitive to monetary policy, while still remaining relatively unaffected by term premia. That said, we also construct instruments based on 3-month, 6-month and 2-year EONIA swaps and do not find a significant difference in our results.

For their high frequency analysis of US monetary policy, Gertler and Karadi (2015) choose a window of 30 minutes around the policy announcement (starting 10 minutes before the Federal Open Market Committee (FOMC) announcement and ending 20 minutes after). The main policy announcement of the FOMC contains a large amount of information about the decision as well as the view of the committee about the state of the economy and expectations of future policy action. This means that within the 30 minute window, the market can fully integrate recent policy changes and adjust the price of the instrument. The procedure of policy releases is somewhat different at the ECB. The release of the monetary policy decision at 13:45 CET only contains a limited amount of information on the latest policy actions. A significant amount of information is disseminated to the market at a later stage, through the press conference and Q&A with the President, starting at 14:30 CET. For this reason, we decided to extend the window for our analysis to cover not only the prime release, but also the press conference. Specifically, we choose a 6-hour window from 13:00 to 19:00 CET.[10]

Figures 1.1 and 1.2 show examples of characteristic movements in the 1-year EONIA swap on ECB meeting days, highlighting the importance of including the Q&A in the high-frequency window. On 5 June 2008, the Governing Council of the ECB decided that policy rates will remain unchanged. As this was in line with market expectations, the 1-year EONIA swap rate did not move much in reaction to the press release at 13:45 CET. During the press conference however, the president expressed concern about increased risks to price stability, setting expectations of rate hikes in the near future. In reaction to this information, the swap rate immediately jumped higher and over the afternoon increased by 27 basis points. This example clearly demonstrates

---

[10]The press conference typically lasts for only one hour, implying that the window could be more narrowly defined, ending, e.g. at 16:00 CET. We chose not to do so due to data availability issues. Specifically, intraday data on swap prices on Bloomberg are available only from January 2008 onwards. In other words, we would have been able to create an instrument only from 2008 using intraday data. For a window from 13:00 to 19:00 CET, however, this problem does not arise as these times correspond to the closing times of the Tokyo and London stock exchanges, respectively. Hence it is possible to obtain end-of-day data, which is available from 2001, and create a *de-facto* intraday window from 13:00 to 19:00 CET. For the subsample of overlapping observations (2008-2016) we tested for the difference in using the window ending with the press conference vs. later the same afternoon and found it to be statistically insignificant.
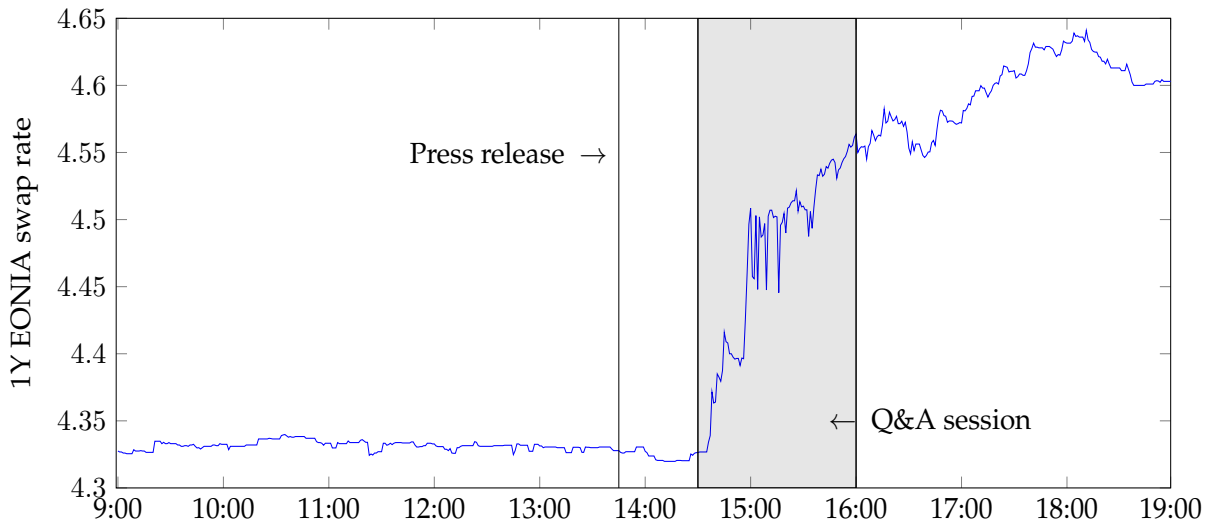
Figure 1.1: 1-year EONIA swap rate on 5 June 2008. Horizontal axis shows Central European Time (CET). Source: Bloomberg, authors' calculations.

that information about ECB policy can to a large degree be contained in the press conference, compared to the policy announcement. An example where both the original announcement, as well as the press conference convey substantial information to market participants is the meeting on 6 October 2011. The press release once again stated that rates would remain unchanged. However, this was not in line with market expectations for a cut and hence created a tightening surprise that led to an immediate increase in the 1-year EONIA swap rate. During the press conference, the then ECB President Jean-Claude Trichet re-emphasised that inflation rates had remained at elevated levels. This in turn pushed market expectations towards tighter monetary policy and caused a further jump in the swap rate. Naturally, there are also examples where the press conference does not convey a significant amount of information to the market, but the above cases highlight the need to include the press release in the high-frequency window.

The above discussion raises the question to which degree the various forms of information dissemination could be used to develop a more differentiated understanding of the nature of policy shocks. Previous contributions have suggested a separation of monetary policy *instrument* shocks from monetary policy *communication* shocks, sometimes also termed *target* and *path* shocks. For the euro area, work on this distinction is being pursued by Jarocinski and Karadi (2018). For the purpose of our paper, we want to use a broad measure of monetary policy shocks that encompasses all forms of surprises, whether they are to the instrument or expectations. That said, we hope that the development of our instrument is informative for future research on the manifold nature of monetary policy shocks within the euro area.

As we estimate a quarterly VAR, we have to turn the surprises on ECB meeting days into quarterly average surprises. In practice, we first calculate the cumulative daily surprise over
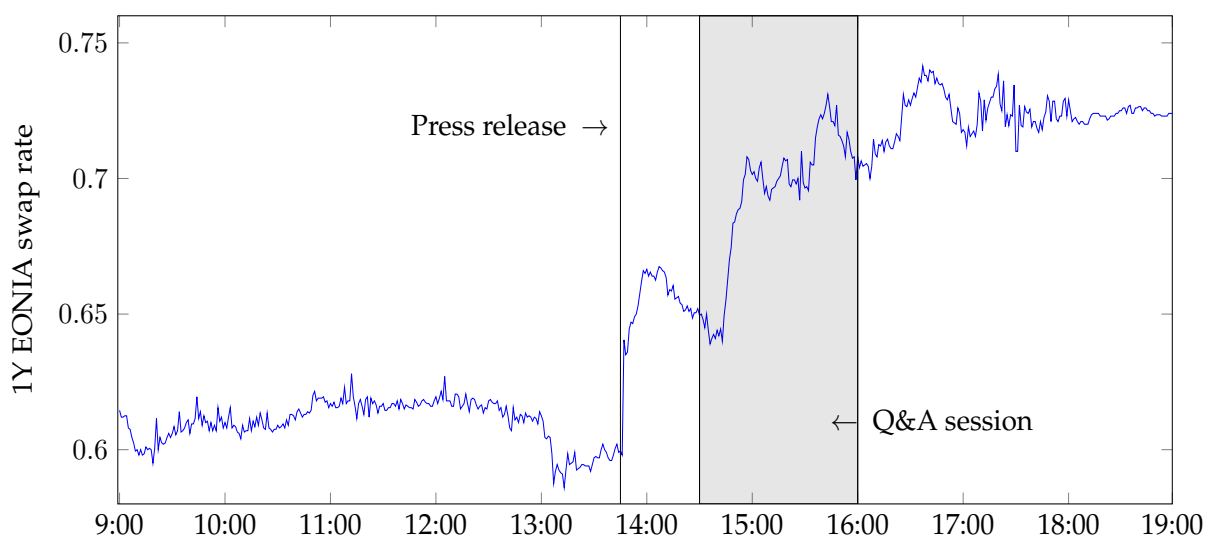
Figure 1.2: 1-year EONIA swap rate on 6 October 2011. Horizontal axis shows Central European Time (CET). Source: Bloomberg, authors' calculations.

the past quarter (93 days) for each day in our sample. In the next step we take the average of this daily cumulative series over each quarter. In doing so, we incorporate the information that some meetings happen early within a quarter while others happen later. Our averaging procedure makes sure that a surprise happening late in the quarter has less influence on the quarterly average than a surprise at the beginning of the quarter.[11]

To get a better understanding of the instrument, we plot its values in Figure 1.3. In particular, we want to point out events that led to particularly large positive or negative values in the instrument to develop an intuition regarding the behaviour of the series. Proceeding chronologically, the earliest of the four largest surprises happened in the fourth quarter of 2001, with a value of -0.15. This data point is driven by the aggressive interest rate cut on 17 September 2001, in response to the 9/11 terrorist attacks.[12] The ECB cut all three interest rates by 50bp leading to a drop in 1-year EONIA swaps of 20bp during our window. Another particularly large negative shock appears in the fourth quarter of 2008. The value of -0.17 is mostly driven by the monetary policy decision on 2 October 2008. Interest rates were kept unchanged on the day, in line with expectations. However, President Trichet highlighted financial market turmoil and weakness in the EA economy during his statement, leading to a large drop in the swap rate between 14:30 and 15:30 CET as markets priced in future cuts to the policy rate. In the following quarter, Q1 2009, our instrument records a particularly high reading of 0.14. This goes back in large part to a contractionary monetary policy surprise during the meeting of 4 December 2008, but also to a surprise during the meeting of 15 January 2009. Interestingly,

---

[11]A similar approach was taking by Gertler and Karadi (2015) to create monthly FOMC surprises.

[12]Note that the surprise actually happened in the third quarter of 2001. However, because our averaging approach takes into account whether a shock appears early or late in a quarter—and consequently, whether it has a larger influence on the current or the next quarter—the policy decision from 17 September 2001 mostly affects our instrument during Q4 2001.
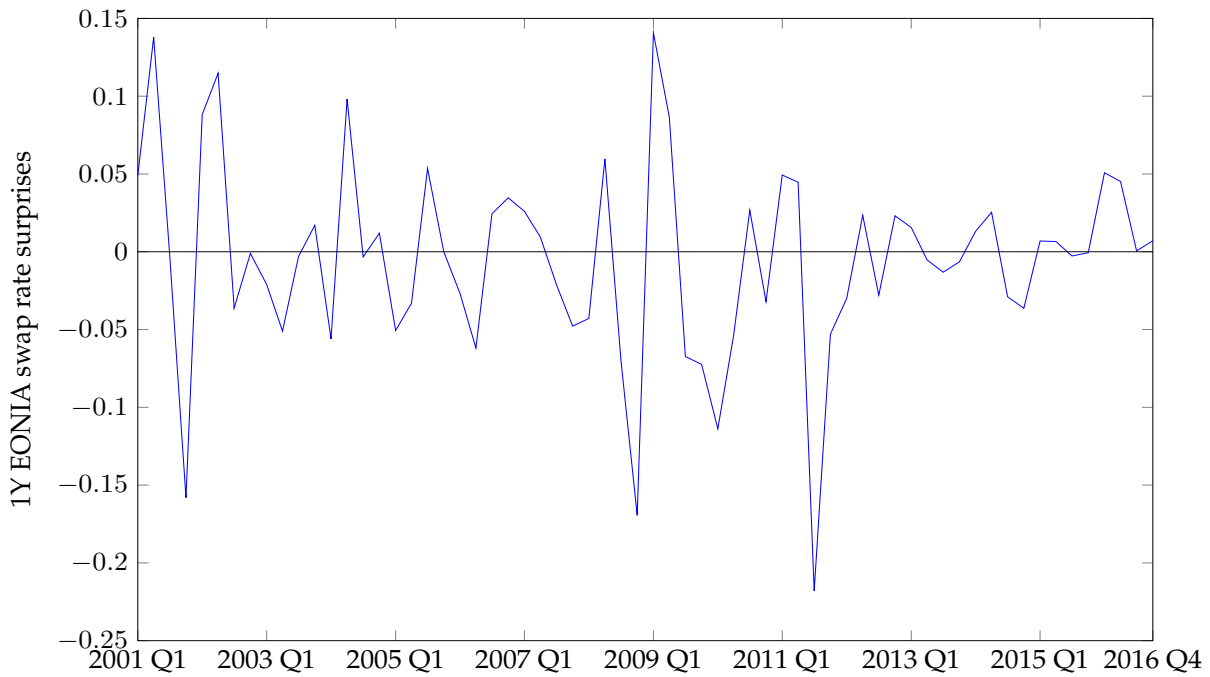
Figure 1.3: Instrument - Quarterly 1-year EONIA swap rate surprises from 2001Q1 to 2016Q4

during both meetings, which happened at the height of the financial crisis, interest rates were cut—by 75bp and 50bp, respectively. While this led to momentarily lower swap rates on both occasions, rhetoric during the press conference led to further increases in the rate. In fact, on both occasions, the President's various dovish and hawkish comments led to the rate moving up and down, but the contractionary sentiment dominated overall. Finally, we investigate the events driving our instrument during Q3 2011. The negative value of -0.22—the largest value in absolute terms during our sample period—mainly goes back to the policy decision on 4 August 2011. After an interest rate hike at the previous meeting, policymakers left interest rates unchanged on the day. As this was in line with expectations, the swap rate did not move at 13:45 CET. During the press conference, however, the ECB announced the decision to conduct a liquidity-providing supplementary longer-term refinancing operation (LTRO), based on observed tensions in financial markets within the euro area. This policy action amounted to a large dovish surprise and 1-year EONIA swaps fell by about 18bp between 14:30 and 15:30 CET.

Finally, we test the strength of our instrument. We do so in a small VAR containing only three variables: output, consumer prices and a policy indicator. The model is specified both at monthly and quarterly frequency and is identified using high-frequency instruments based on 3, 6 and 12-month EONIA swaps. We report further details and all results in Appendix B, but note here that in our baseline specification the instrument is strong, with a first-stage F-test statistic of 19.45. This confirms the relevance of our external instrument.

14

1.2.4 Data and Estimation

Our data set consists of quarterly observations from 1999 Q4 to 2016 Q4 on 90 area-wide measures such as prices, output, investment, employment and housing, as well as 342 individual country time series for the 11 early adopters of the Euro: Austria, Belgium, Finland, France, Germany, Ireland, Italy, Luxembourg, the Netherlands, Portugal and Spain. The vintage of the data is June 2017. Appendix C lists all data series with detailed descriptions and notes on the completeness and length of the individual series.

All data series are transformed to induce stationarity. Depending on the nature of the data, this was done either by taking the first difference in logs or levels. Details on transformations can also be found in Appendix C. As we lose one observation by differencing, our working dataset starts in 2000 Q1.

Principal component analysis is sensitive to double-counting[13] and we consequently only use a subset of our data for factor extraction. In practice, we avoid double-counting along two dimensions. Firstly, we do not include euro area aggregates for indicators where we have included all individual country series. Secondly, we do not include category aggregates, such as GDP, when we have included its components, such as the components of GDP. Where possible, we avoid using high-level aggregate series altogether and instead include disaggregate series. In total, we use 179 series for factor extraction.

Table 1.1: Determining the number of common factors: Onatski (2009) test. The Table shows p-values of the null of q_0 common shocks against $r_0 < r \leq r_1$ common shocks.

r_0 **vs** $r_0 < r \leq r_1$	1	2	3	4	5	6	7
0	0.727	0.089	0.122	0.153	0.18	0.209	0.232
1	0	0.05	0.089	0.122	0.153	0.18	0.209
2	0	0	0.521	0.414	0.539	0.632	0.705
3	0	0	0	0.229	0.414	0.539	0.632
4	0	0	0	0	0.794	0.595	0.746
5	0	0	0	0	0	0.336	0.595
6	0	0	0	0	0	0	0.561

We rely on a number of specific tests and information criteria to determine the number of common factors r. Specifically, we estimate them by means of the test proposed by Onatski (2009), which suggests $r \in 2, 3$ (Table 1.1), the eigenvalue difference method proposed by Onatski (2010) suggesting $r = 2$, the criterion by Bai and Ng (2002) suggesting $r = 5$, and the bi-cross-validation method proposed by Owen and Wang (2015)[14] suggesting $r = 8$. We choose as our baseline specification $r = 5$, that is, the average of these results. Figure 13 in Appendix A shows

---

[13]see e.g. Stock and Watson (2012).
[14]see Figure 12 in Appendix A.

Table 1.2: R-squared for regression of data series on five principal components. *Germany, France, Italy, Spain, Netherlands. **Belgium, Austria, Ireland, Finland, Portugal, Luxembourg.

	EA aggregate	Average across individual country series	Average across large* countries	Average across small** countries
Gross Domestic Product	0.85	0.56	0.70	0.45
Harmonised Index of Consumer Prices	0.81	0.64	0.71	0.59
House Prices	0.71	0.46	0.52	0.40
Exports	0.76	0.54	0.49	0.58
Imports	0.75	0.58	0.45	0.69
Government Spending	0.18	0.68	0.77	0.59
Gross Fixed Capital Formation	0.76	0.33	0.51	0.19
Consumption	0.61	0.30	0.34	0.27
Unemployment	0.72	0.51	0.68	0.36
Long-term Rates	0.99	0.98	0.98	0.98
Rents	0.41	0.35	0.32	0.38
Share Prices	0.65	0.58	0.59	0.57
Producer Prices in Industry	0.87	-	-	-
Wages	0.75	-	-	-
Employment	0.74	-	-	-
GER 2Y yield	0.98	-	-	-
Cost of Borrowing indicator	0.91	-	-	-
EONIA	0.99	-	-	-
Nominal Effective Exchange Rate	0.12	-	-	-

the variance of the data explained by each additional factor. Five factors account for 80% of the total data variance.[15]

On the basis of Akaike and Bayes Information Criteria we include one lag for the baseline of the DFM.

To get a better understanding of how well the extracted factors characterise the data, Table 1.2 shows the variation in the data explained by the five factors. The second column shows the fraction of explained variation for a selection of aggregate area-wide series. The third column shows the corresponding average across series from individual member countries. In particular, two observations stand out. Firstly, the variation in most aggregate series is remarkably well explained by the five factors. With a few exceptions, notably the exchange rate, the R-squared ranges between 70% and 99%. Secondly, despite the granularity of the individual country series, the factors on average still explain more than half of all variation. In some cases, such as HICP inflation, government spending and, most notably, long-term interest rates, they explain

---

[15]As can be seen in Figure 13, the bulk of the variance in the data is explained by the first two factors. In line with this observation and the test results from Onatski (2009) and (2010), we re-estimate the DFM with only two factors. We find that all main results of the 5-factor model hold. While the smaller amount of factors allow for greater precision, the larger amount of factors gives us more explanatory power for the observable series. We prefer the latter effect over the former and hence select 5 factors for our baseline specification.

considerably more. Columns 4 and 5 show the same information as column 3, but differentiate between the size of the countries. In particular, we separate the 5 countries in our sample with the largest economies (by nominal GDP) from the 6 countries with the smallest economies. As expected, the factors pick up information from the large economies to a much greater extent than for smaller economies. With the exception of exports, imports and rents, data from larger economies is consistently explained better by the factors. This difference is particularly strong for GDP (70% vs. 45%) and unemployment (68% vs. 36%). As concrete examples of the above, Figure 17 in Appendix E plots fitted series on the basis of the 5 extracted factors against actual (transformed) series for GDP and HICP in the euro area, Germany and Luxembourg.

1.3 Empirical Results

This section gives an overview of our empirical findings, starting at the euro area aggregate level and subsequently exploring results on the country level.

1.3.1 Euro-wide Dynamic Effects of Monetary Policy

We start our description of the results with an overview of a selection of aggregate series across the euro area. Figure 1.4 shows percentage responses to a contractionary monetary policy shock of 25 basis points (bp). As discussed in Section 1.2.3, the external instrument approach identifies the shock only up to sign and scale. Using the response of EONIA as a policy indicator, we scale the system to a 25bp contraction in EONIA. The shaded area around the point estimates signify confidence intervals of one standard deviation, obtained from a wild bootstrapping procedure with a simple (Rademacher) distribution. Given a strong instrument, the confidence intervals obtained under this approach are valid despite the presence of heterogeneity. Because both stages of the regression are incorporated in the bootstrapping procedure, the error from the external instrument regression is accounted for. A similar approach has been followed by Mertens and Ravn (2013) and Gertler and Karadi (2015).
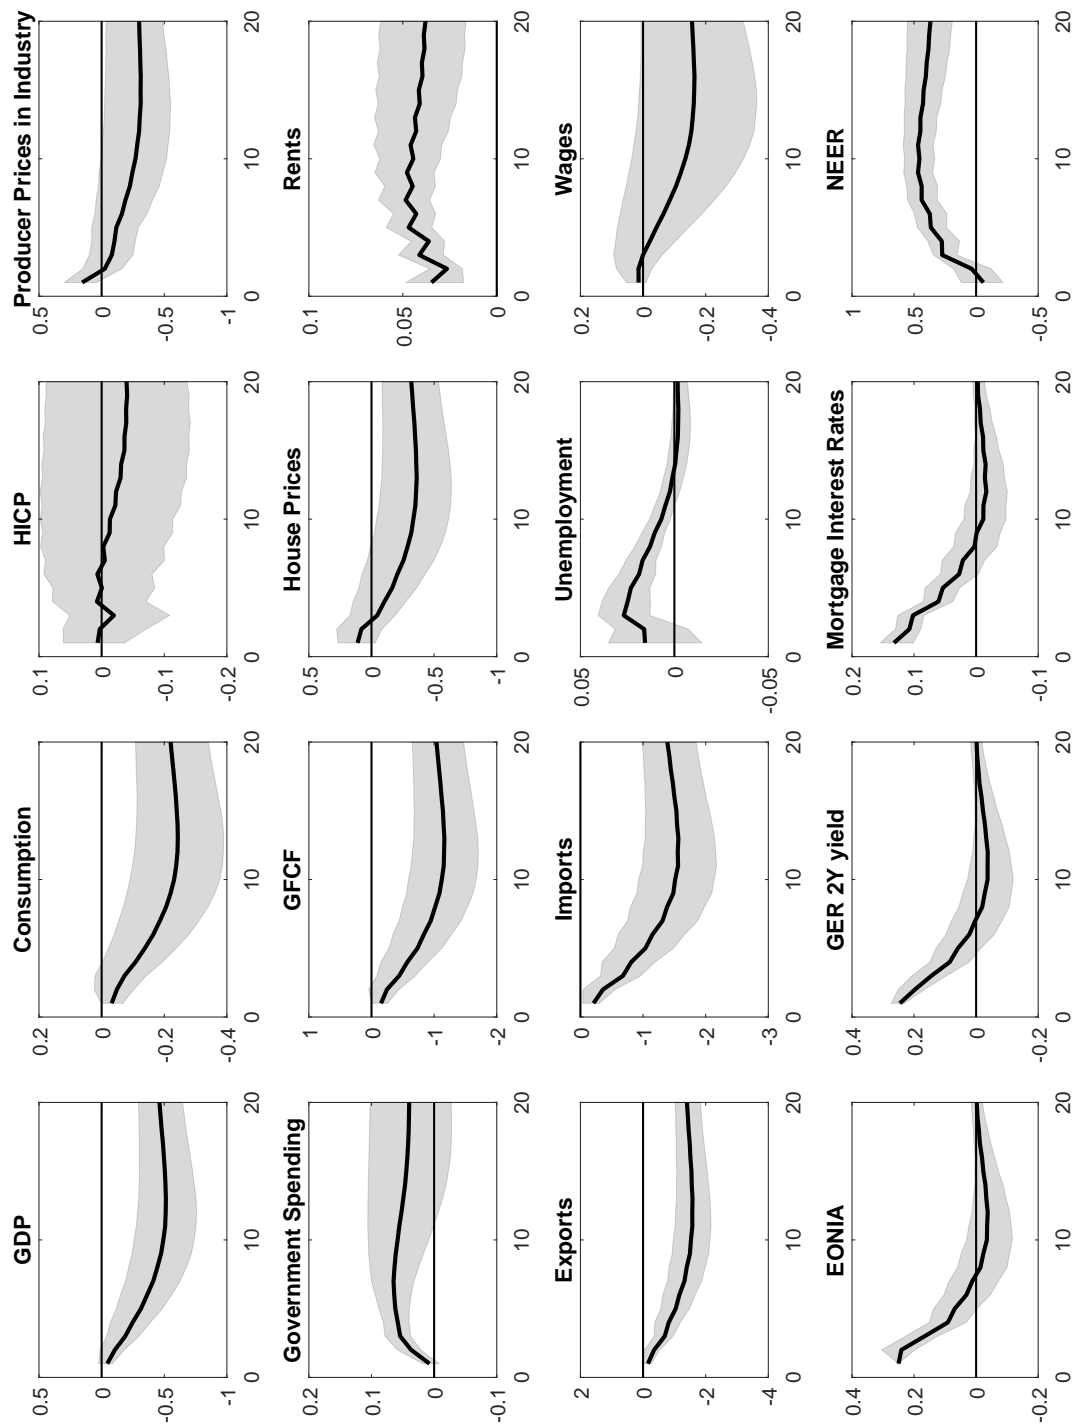
Figure 1.4: Percentage responses of selected euro-wide variables to a 25bp contractionary policy shock. Note: Confidence intervals are obtained from a wild bootstrap procedure with a simple (Rademacher) distribution.

Notably, our results do not suffer from the prize puzzle—the occurrence of rising prices in reaction to a contractionary monetary policy shock. In fact, while the harmonised index of consumer prices (HICP) does not have any significant reaction, our producer prices fall significantly, in line with economic theory. Given the longstanding struggle of VAR-type models to get rid of the price puzzle, we interpret these findings as an indication of the ability of the model to accurately characterise economic dynamics. In particular, we attribute the non-existence of the price puzzle to the combination of correctly capturing information about prices in the economy (via the DFM) and precisely identifying monetary policy shocks (via the high frequency instrument).[16] The remainder of the series in Figure 1.4 also behave as suggested by theory. GDP contracts overall, as do all components with the exception of Government Spending, which increases in reaction to a contractionary shock. In line with theory, investment (GFCF) is a lot more volatile than consumption, as are imports and exports. The reaction of the German 2-year sovereign yield closely follows EONIA. The aggregate indicator for mortgage interest rates in the euro area as compiled by the ECB also rises in reaction to a shock, but displays imperfect pass-through as a significant number of mortgages are characterised by fixed rates that do not adapt to changes in policy. In the labour market, unemployment rises, while wages fall. Interestingly, the reaction in wages is not significant, hinting at a large degree of nominal wage stickiness. In the housing market, house prices fall significantly after a contraction, following economic theory that higher policy rates make mortgages more expensive and consequently suppress demand for houses. Rents, on the other hand, increase in reaction to a shock. Recent research (see e.g. Duarte and Dias, 2016) suggests that a worsening of conditions in the mortgage market leads agents to substitute house purchase with renting, thus exerting pressure on rental prices. Motivated by this result, Section 1.4 will take a closer look at the housing markets in the euro area and explore potential avenues for connecting the results to economic theory.

1.3.2 Country-Level Dynamic Effects of Monetary Policy

Moving on to results at the country level, we start to uncover the full potential of the DFM when it comes to providing results for a large number of series. Of the 342 individual country series in our data set, we have selected a representative sub-sample for Figures 1.5-1.7. In particular, this section takes a closer look at the responses of GDP, the components of GDP, interest rates, equities, housing prices and housing rents. We point out, however, that the model produces impulse response functions for all series in our sample.

Figure 1.5 shows responses for real GDP and HICP across the 11 euro area countries in our sample. While we omitted error bands for ease of presentation, it is noteworthy that reactions of

---

[16]We also applied the FAVAR approach proposed by Bernanke et al. (2005) using EONIA as the only observable factor and found that the price puzzle was still present
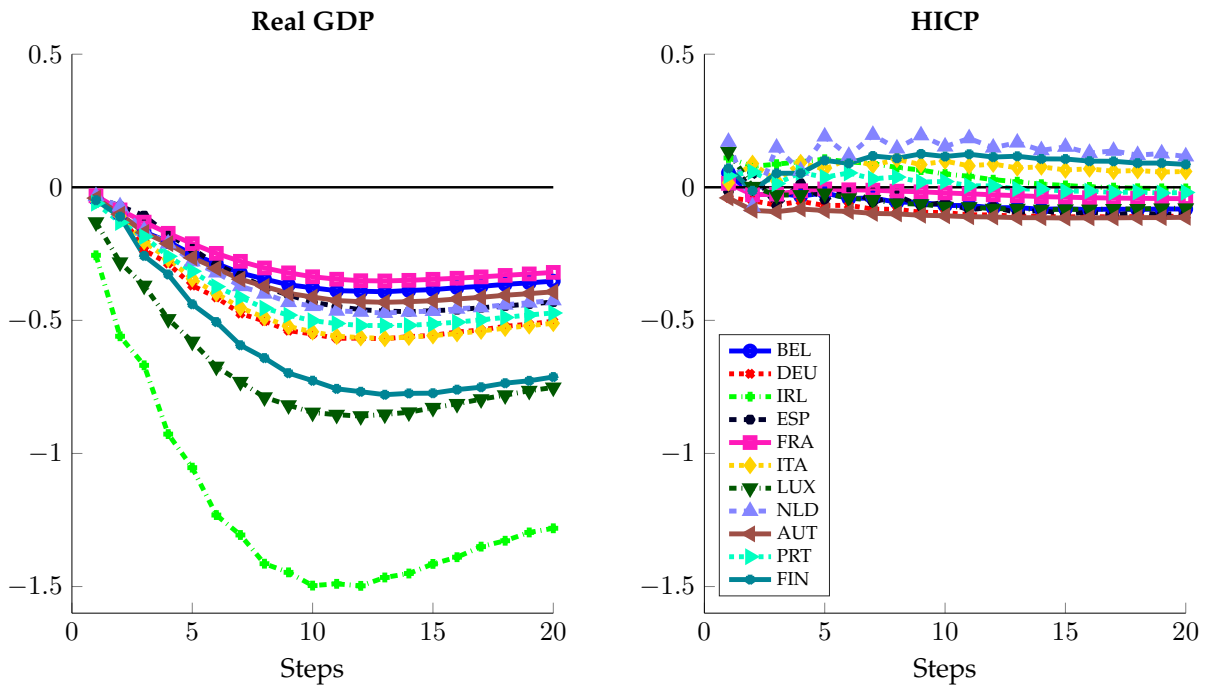
Figure 1.5: Percentage responses of real GDP to a 25bp contractionary policy shock across euro area member countries.

GDP across countries are significantly heterogeneous. At one end of the spectrum, the reaction of Irish GDP clearly differs from the five countries with the weakest reaction. That said, even the reactions of Finland and Luxembourg are statistically different from France and Spain, having non-overlapping confidence intervals from the 10th step onward. This heterogeneity is in itself noteworthy, but also raises the question which parts of the economy are particularly prone to asymmetric reactions.

For a first pass at this question, Figure 1.6 contains the reactions of the components of GDP. A look at the IRFs offers two main conclusions. Firstly, some series, such as private consumption and gross fixed capital formation tend to move in the same direction, or with similar patterns, despite heterogeneity across countries. This compares to series such as government spending and net exports, which in some instances even move in opposite directions. In part, these differences in the general nature of responses can be explained by the determinants of the individual series. Government spending, for example, is notoriously idiosyncratic, varying in degrees of pro- and countercyclicality both across countries as well as within a country over time.

Secondly, we observe that even among series where responses across countries move in the same direction, large degrees of heterogeneity exist. In particular, we point out the disparity in reactions of private consumption. While German private consumption drops by a maximum of about 0.02 percentage points, the drop in Ireland is more than 20 times as large at 0.4 percentage points. Aside from Ireland, which could be classified an outlier, countries such as Italy, Finland, Spain and Portugal exhibit drops in consumption that are roughly 10 times the size of

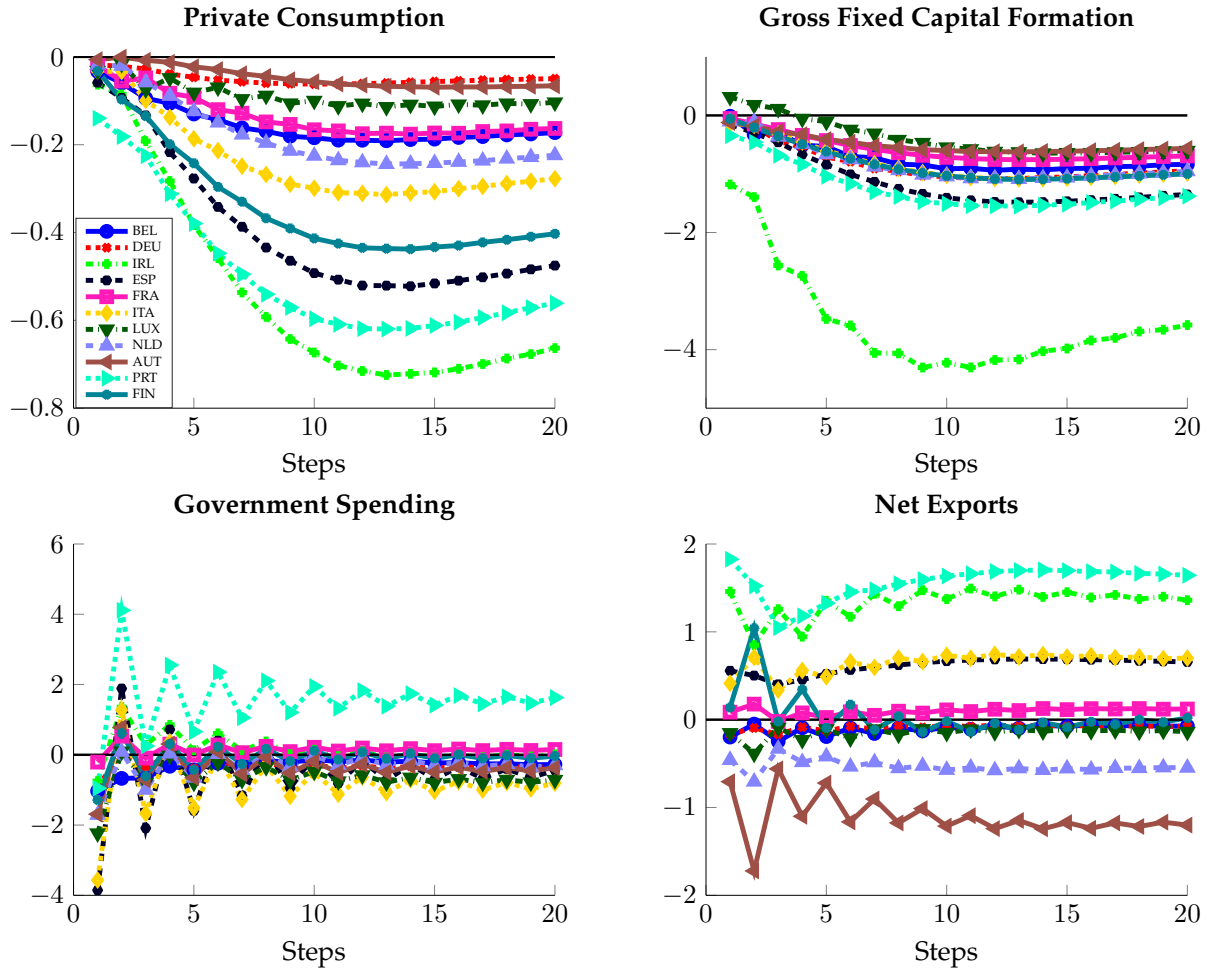**Figure 1.6:** Percentage responses of GDP components to a 25bp contractionary policy shock across euro area member countries.
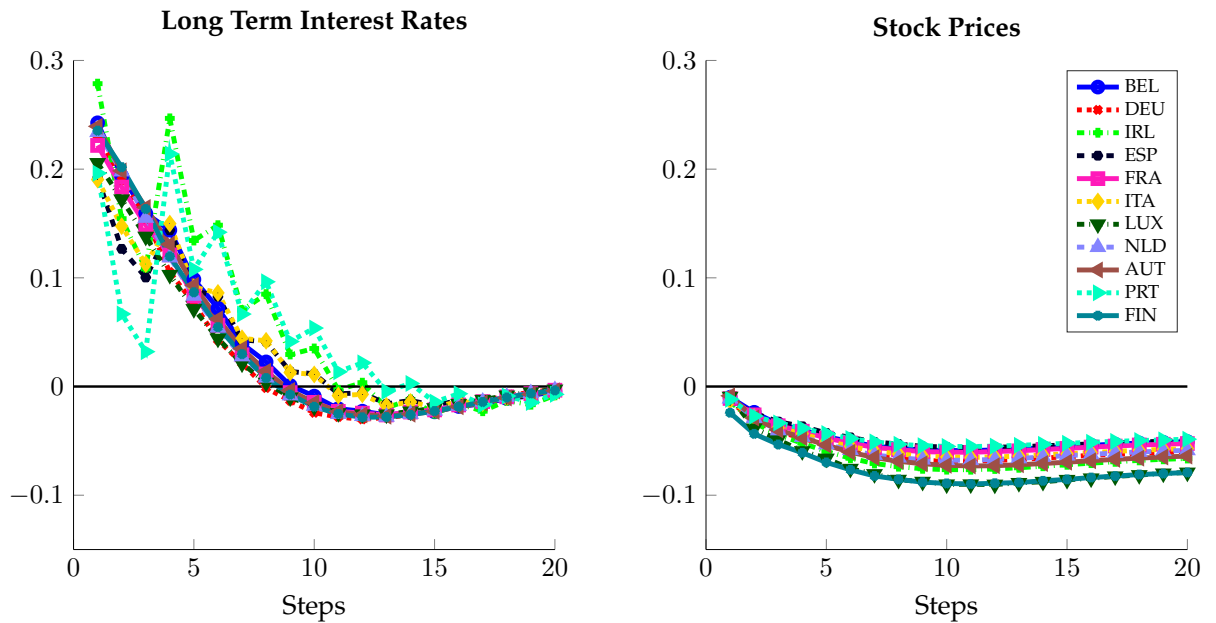

Figure 1.7: Percentage responses of long-term interest rates and local equity indices to a 25bp contractionary policy shock across euro area member countries. Long-term interest rates are defined in accordance with OECD methodology, conforming to government bonds of (in most cases) 10 year maturity.

Housing Prices   **Housing Rents**
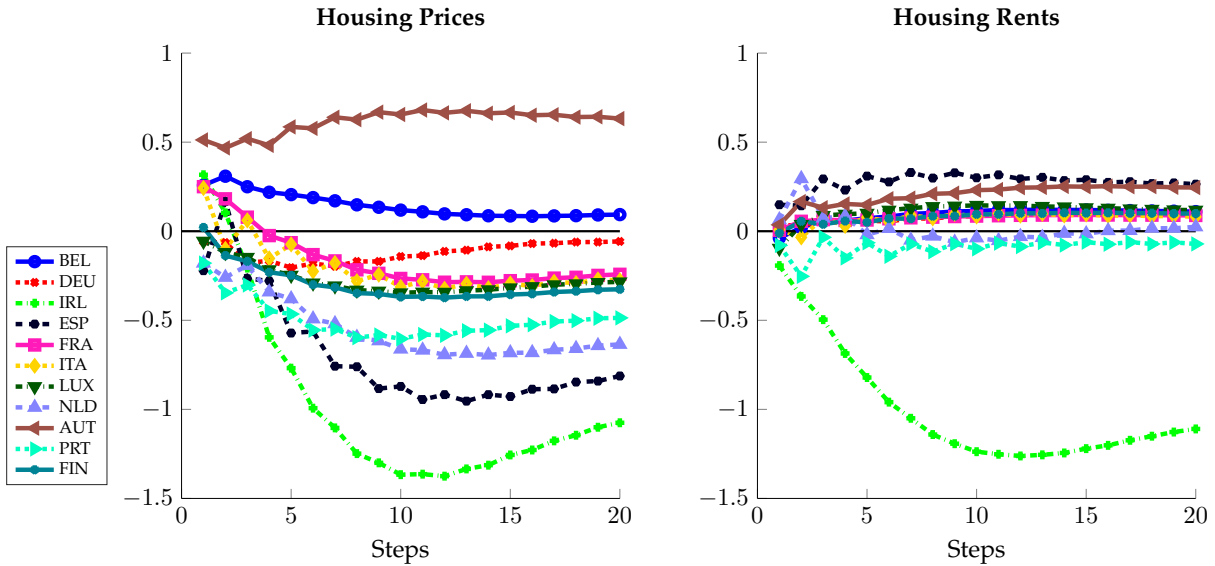
Figure 1.8: Percentage responses of house prices and rents to a 25bp contractionary monetary policy shock across euro area member countries.

the reaction in Germany. One of the core questions we ask in Section 1.4 is what may be the cause of this degree of heterogeneity.

Taking a closer look at responses of other variables of interest, we find that the degree of heterogeneity in responses seems to be closely (and inversely) related to the state of convergence in a particular market across the euro area. In particular, financial markets have seen a large degree of convergence,[17] which is reflected in the reaction of interest rates and stock prices across countries. Figure 1.7 shows that while the immediate impact of a policy shock on long-term interest rates is not uniform across countries, their reaction over later periods is almost identical. Similarly, a look at the responses of local equity indices, displayed in the same figure, reveals a strong degree of homogeneity across equity markets. While stock prices show more differentiated responses than long-term interest rates, the confidence intervals around stock price IRFs are mostly overlapping.

Among the markets that have seen only very little or no convergence in institutional characteristics are the labour and housing markets. It is the latter, seen in Figure 1.8, that we investigate in more detail in Section 1.4, thus attempting to clarify the role of institutional characteristics in shaping channels of the monetary transmission mechanism.[18]

In the following step, we propose a more rigorous approach to test heterogeneity among

---

[17]see e.g. ECB (2017).

[18]In Appendix F, we present an alternative way of looking at IRFs across countries to better understand the statistical significance of our results. Figures 18 and 19 plot the highest and lowest responses, as well as EA IRFs for various series with their respective confidence intervals. Figure 18 contains these responses for real variables: GDP, private consumption and unemployment. Figure 19 shows responses for price-related series: interest rates, HICP and stock prices. Comparing the two groups, we notice that the highest and lowest responses for none of the real variables are overlapping. In contrast, IRFs are overlapping for most parts of the price-related series, with only stock prices diverging around the 10th step.

Table 1.3: Coefficient of variation of the cross-country responses to a 25bp monetary policy shock.

Variable	Coefficient of Variation	Lower Bound	Upper Bound
On Impact			
GDP	1.45	0.70	4.00
Private Consumption	1.19	1.01	2.52
Unemployment Rates	7.16	2.83	25.02
Housing Prices	2.03	1.51	4.57
Housing Rents	3.81	3.15	6.96
HICP	3.24	0.99	13.25
Long-term Interest Rates	0.21	0.14	0.53
Stock Prices	0.37	0.21	0.65
At the 20th Step			
GDP	0.64	0.47	0.95
Private Consumption	1.02	0.99	1.11
Unemployment Rates	1.24	0.94	4.22
Housing Prices	1.08	0.84	2.02
Housing Rents	2.41	1.13	8.20
HICP	1.25	0.62	4.05
Long-term Interest Rates	0.46	0.17	1.87
Stock Prices	0.21	0.19	0.26

responses. For each set of responses, we calculate the coefficient of variation, i.e. the standard deviation of responses (among countries) with respect to the EA response of the same variable. To make this measure comparable across different types of series, we normalise it by the size of the EA response. In doing so, we create a numerical measure for the dispersion of impulse responses that allows for intuitive and meaningful comparison between series. Table 1.3 presents the coefficients of variation for a selection of variables, both on impact, as well as at the 20th step. Moreover, the table lists a lower and upper bound for the coefficients of variation, which we obtain from including the calculation in our bootstrapping procedure. As can be clearly seen, long-term interest rates and stock prices have a much smaller coefficient of variation than the other presented variables. At the 20th step, GDP is also markedly less heterogeneous than other variables such as private consumption.

As some of the intervals around coefficients of variation are overlapping, we also bootstrap pair-wise differences in the coefficient of variation. The results, presented in Table 1.4, mostly confirm earlier observations. Reactions of long-term interest rates (LTINT) and stock prices (SP) are significantly less dispersed than all other variables. Moreover, at the 20th step, GDP has a significantly lower coefficient of variation than private consumption (PCON), unemployment (U), real house prices (RHPI) and real rents (RREN). A few additional pairwise tests also show significance.

Overall, the findings confirm our earlier statement: The degree of heterogeneity in responses is lower in financial variables, such as interest rates and stock prices, as well as output, while it is larger in consumption, consumer prices and variables related to the labour and housing markets. Importantly, these results open up scope for potentially beneficial policy intervention. Further institutional convergence might increase the efficiency and precision of monetary policy, reducing unintended reactions in countries that would otherwise behave in an idiosyncratic manner. That said, a much deeper understanding of the mechanisms at play is necessary to justify policy intervention in the first place. By means of a case study, the following section takes the first step in the direction of improving our understanding for the example of the housing market.

Table 1.4: Bootstrapped pair-wise differences in the coefficient of variation of the cross-country responses to a 25bp monetary policy shock. * marks differences in variation that are significant at the 68% confidence level. The inference is drawn from a bootstrap procedure.

	GDP	HICP	LTINT	SP	PCON	U	RHPI	RREN
On Impact								
GDP	0	-0.99	1.20*	1.06*	0.16	-5.42*	-0.84	-2.15
HICP	0.99	0	3.02*	2.85*	1.69	-3.81	0.66	-0.41
LTINT	-1.20*	-3.02*	0	-0.13	-0.90*	-6.66*	-1.79*	-3.43*
SP	-1.06*	-2.85*	0.13	0	-0.84*	-6.84*	-1.60*	-3.32*
PCON	-0.16	-1.69	0.90*	0.84*	0	-5.20*	-0.75	-2.48*
U	5.42*	3.81	6.66*	6.84*	5.20*	0	5.02	3.46
RHPI	0.84	-0.66	1.79*	1.60*	0.75	-5.02	0	-1.51
RREN	2.15	0.41	3.43*	3.32*	2.48*	-3.46	1.51	0
At the 20th Step								
GDP	0	-0.55	0.21	0.45*	-0.39*	-0.59*	-0.43*	-1.74*
HICP	0.55	0	0.64	1.02*	0.19	-0.18	-0.16	-0.93*
LTINT	-0.21	-0.64	0	0.24	-0.60	-0.99*	-0.62	-1.65*
SP	-0.45*	-1.02*	-0.24	0	-0.80*	-1.04*	-0.85*	-2.17*
PCON	0.39*	-0.19	0.60	0.80*	0	-0.20	0.00	-1.38*
U	0.59*	0.18	0.99	1.04*	0.20	0	0.20	-0.84
RHPI	0.43*	0.16	0.62	0.85*	0.00	-0.20	0	-0.66
RREN	1.74*	0.93*	1.65*	2.17*	1.38*	0.84	0.66	0

1.4 Case Study of Heterogeneous Monetary Policy Transmission in Euro Area Housing Markets

The recent literature has emphasized how the transmission of monetary policy operates through a "housing transmission channel".[19] A focus on this channel is commonly motivated by noting that for most households, the single most important item on the asset side of their balance sheet

---

[19]Among the various contributions, Iacoviello (2005) and Kaplan et al. (2016) call attention to credit and liquidity constraints, while the role of rents is studied in Duarte and Dias (2016).
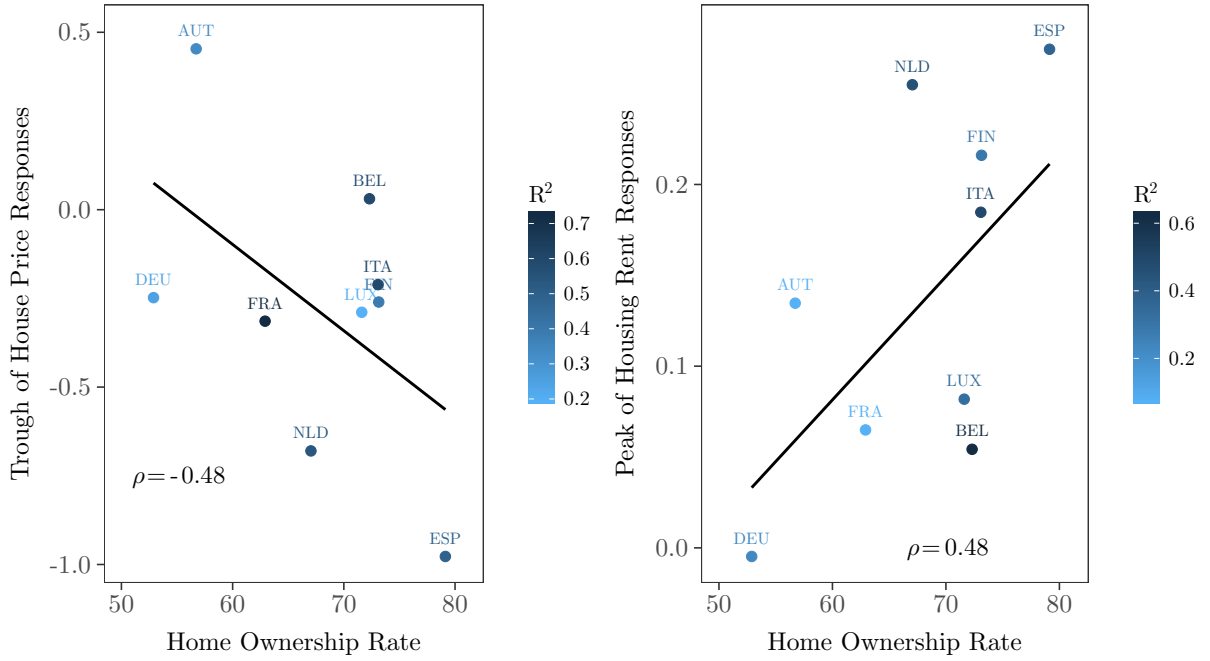
Figure 1.9: Correlation of house price response troughs and rent response peaks with home ownership rates. Figures exclude Ireland and Portugal.

is their home. In the form of the mortgage, it typically corresponds to a household's largest liability. In this section, we employ our dynamic factor model for the euro area to investigate the housing channel in more detail. Concretely, we will make use of the European setting to explore how various details of the housing channel shape the transmission of monetary shocks.

A remarkable feature of European housing markets are their substantial differences in institutional characteristics. Mortgage markets differ markedly in the relative share of fixed versus flexible rate contracts and maximum loan-to-value ratios, rental markets are subject to different regimes and controls, and property taxation is very heterogeneous, to name but a few aspects— see Osborne (2005), Andrews et al. (2011) and Westig and Bertalot (2016) for a comprehensive overview. Indeed, the effect of institutional characteristics on the transmission of monetary policy has previously been explored by Calza et al. (2013), who however focus mostly on housing finance.

Our point of departure is the idea that the transmission of monetary policy should be stronger in an environment where rents and house prices move more in response to a shock. This—as we explain below—may occur where home ownership rates are high.

To bring this idea to the data, we use our dynamic factor model to trace out the transmission of a monetary policy shock through house prices and rents, as well as mortgage rates and cost of finance, with the goal of shedding light on each constituent component of the housing channel.

Figure 1.8 shows impulse response functions for house prices and rents across the euro area. As anticipated, house prices and rents feature a strong degree of heterogeneity, both in a qualita-

tive as well as quantitative sense. In particular, we note that while the majority of house prices fall in response to a contractionary monetary policy shock, rents in fact *rise* in most countries. This is also reflected in the aggregate responses of euro area house prices and rents as seen in Figure 1.4. This dichotomy may seem counterintuitive when looking at rents and house prices from an asset pricing perspective: House prices should reflect the discounted sum of expected future rents, leading to parallel movements in both. While an asset pricing approach is helpful in explaining price developments when housing takes on the role of an investment good, it neglects particularities of housing as a consumption good. When an interest rate shock hits the housing market, mortgages become more expensive, which consequently leads to a fall in the demand for houses and hence their price. As highlighted in the recent literature,[20] when faced with more expensive mortgages, home buyers, at the margin, not only scale down the size of their mortgage (and house), but also substitute buying a house with renting. In doing so, demand for rental properties goes up, leading to an increase in rents. In line with our hypothesis above, the ability of a rental market to absorb the mass of agents that switch from buying to renting may be closely linked to the size of the rental market or, inversely, to the home ownership rate of a country. To get a better understanding of the importance of the home ownership rate in the transmission channel, in Figure 1.9 we plot the maximum response of each country's house price and rent IRF against the home ownership rate. Seeing that the responses of Ireland and Portugal are highly idiosyncratic, we treat them as outliers at this stage and do not include them in the scatter plots. Note that the correlation of consumption to rents is weakened by the inclusion of Ireland and Portugal, while all other correlations are strengthened by their inclusion.

Before analysing the scatter plots in detail, a note of caution regarding their interpretation is in order. With only 9 data points, our plots are meant to motivate and inform a structured approach to the monetary transmission channel—to be pursued in future research—rather than uncovering statistically significant relationships.

The right panel of Figure 1.9 plots the peak of rent responses against home ownership rates. The correlation is positive: The higher the home ownership rate, the larger the increase in rents. As indicated above, this finding is in line with the idea that housing is not only an asset, but also a consumption good. To gain insight on how this matters for monetary policy, observe that, if a shock reduces the demand for housing, so that households switch to renting, the vacancy rate in the housing stock for sale temporarily increases—since conversions between homes for sale and homes for renting are limited. Correspondingly, the pressure on the rental market builds up. Indeed, in countries where the rental market is insufficiently deep, as may be the case when the market is small (i.e., the ownership rate is high), everything else equal, new entrants will exert a greater upward pressure on prices. In the euro area, there are strong cross-border differences

---

[20]See Duarte and Dias (2016).
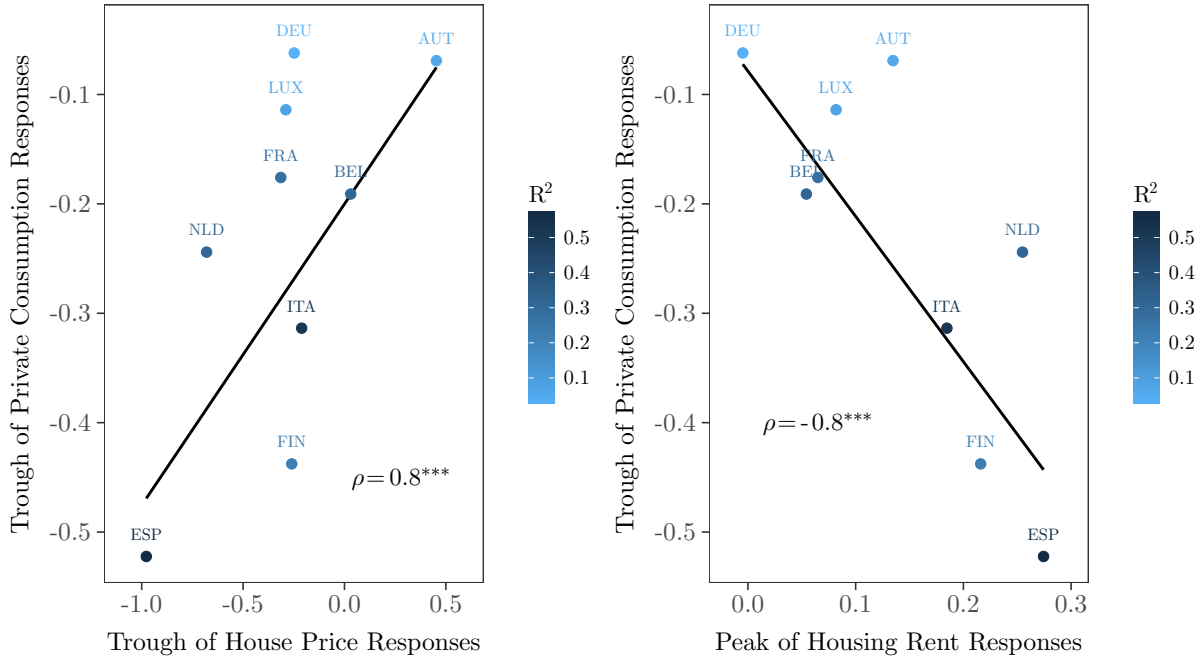
Figure 1.10: Correlation of house price response troughs and rent response peaks with private consumption response troughs. Figures exclude Ireland and Portugal.

in this respect. In Germany, a deep rental market can easily absorb substantial movements at the extensive margin, with agents switching from buying to renting or vice-versa, without experiencing significant variations in rents. At the other extreme, Spain, Finland or Italy, where rental markets are relatively small, reflecting a number of institutional constraints, new entrants can lead to substantial changes in rents.

What remains to be shown is that the rental market actually forms a significant component of the housing channel. To address this issue, we plot the trough (minimum) of house price responses and peak (maximum) of rent responses against the response in private consumption. As can be seen on the two panels, consumption is strongly linked to both changes in house prices as well as changes in rents. The link between house prices and consumption points to a strong wealth effect of monetary policy, stressed by Mian et al. (2013) among others. The larger the drop in house prices, the stronger the direct impact on households' balance sheets, leading to a cut in consumption. Looking at rents, on the other hand, we see a strong negative relationship. Households who pay rent are, on average, less wealthy than households who receive rents, i.e. owners of rental properties. Given a marginal propensity to consume that is decreasing in wealth, this implies that, after an increase in rents, renters as a group cut consumption by a larger amount than landlords increase their consumption, leading to a negative demand shock overall. Moreover, it is plausible that, after a contractionary policy shock (associated with an increase in rents), more renters may become liquidity and credit constrained, causing them to reduce consumption more sharply than implied by any temporary drop in income. In summary,
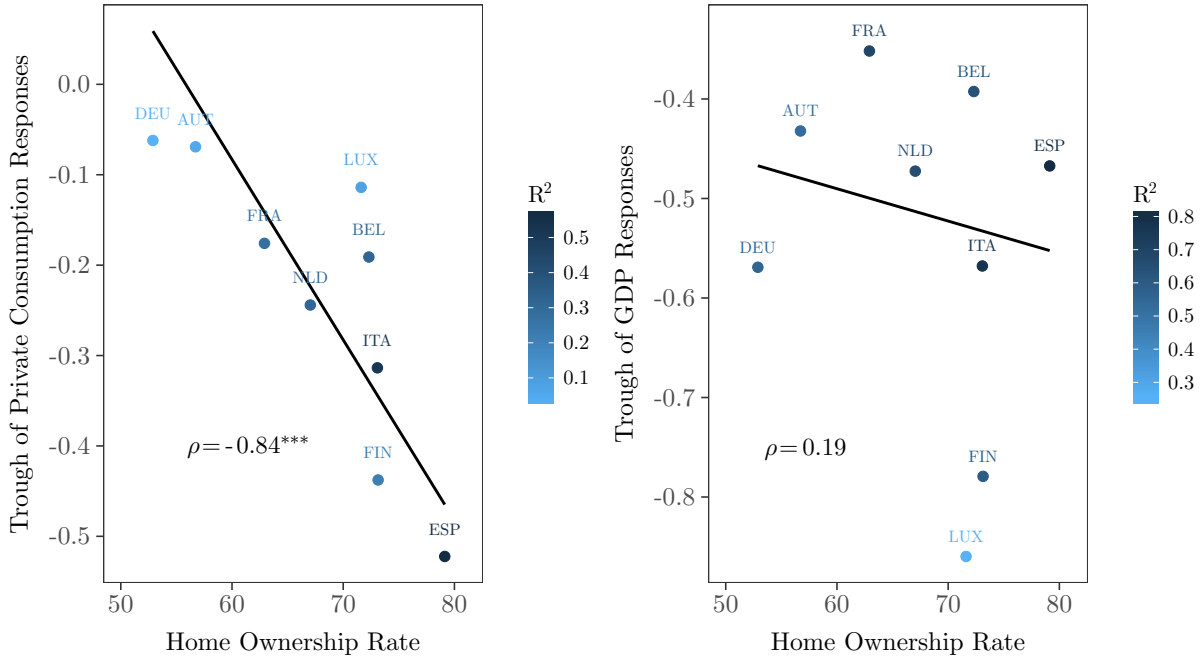
Figure 1.11: Correlation of private consumption and GDP response troughs with home ownership rates. Figures exclude Ireland and Portugal.

our results suggests that the housing channel should be investigated in its multiple components, including house prices and rents. The importance of these components in turn depends on institutional characteristics, such as the home ownership rate, which is plausibly related to housing finance.

Having traced the effect of the home-ownership rate to consumption through house prices and rents, as a last step we investigate whether we can make out a direct relationship between the home ownership rate and changes in consumption. Figure 1.10 plots the trough of private consumption responses against home ownership rates. The figure uncovers a surprisingly clear correlation. This result foreshadows large potential benefit from a systematic analysis of institutional characteristics in relation to the monetary transmission mechanism, in particular in a heterogeneous environment such as the EA. At the same time, it lends support to models stressing the housing channel.

1.5 Conclusion

Using a dynamic factor model with high frequency identification, this paper investigated the heterogeneous effects of monetary policy across the euro area. In doing so, we contribute to the literature by creating an external instrument for monetary policy identification and, by means of a case study on housing, presented a novel way of uncovering heterogeneity in the monetary transmission mechanism. The analysis has produced three main results.

Monetary policy transmission in the euro area appears to be persistently heterogenous across member countries. In this paper, we provided evidence consistent with the idea that the degree of heterogeneity is inversely related to the degree of cross-border institutional convergence. While country-level financial variables and output react fairly similarly to the same monetary policy shock, variables naturally related to markets that have seen little convergence, such as housing and labour markets, react in significantly asymmetric ways.

We elaborate on this point with a case study of European housing markets. We show that differences in the home ownership rate—an indicator reflecting many dimensions in which national housing markets differ from each other—is strongly correlated with the strength of monetary policy transmission across countries. Moreover, we show the importance of looking at the different components of the housing channel, including rents, in addition to house prices. Indeed, our analysis shows that, in most countries, house prices and rents respond to a contractionary policy shock in different directions, yet both contribute to a fall in consumption.

Our results point to a number of promising directions for future research. Firstly, once data availability improves, it would be highly interesting to break down our external instrument and uncover not only shocks to general monetary policy, but also to sub-components, such as monetary policy communications shocks or monetary policy instrument shocks. Secondly, our case study on the housing market is only of exploratory nature. In addition to the home ownership rate, many other institutional characteristics merit a closer look: loan-to-value ratios, aggregate mortgage debt to GDP and ease of credit, to name just a few. Finally, markets other than the housing market, such as the labour market, can and should be examined in a similar fashion.

Bibliography

Andrews, D., Caldera, A., and Johansson, A. (2011). Housing markets and structural policies in OECD countries. OECD Economics Department Working Papers 836, OECD Publishing.

Angeloni, I., Kashyap, A., and Mojon, B., editors (2003). *Monetary Policy Transmission in the Euro Area: A Study by the Eurosystem Monetary Transmission Network*. Cambridge University Press.

Bagliano, F. C. and Favero, C. A. (1999). Information from financial markets and VAR measures of monetary policy. *European Economic Review*, 43(4):825 – 837.

Bai, J. and Ng, S. (2002). Determining the number of factors in approximate factor models. *Econometrica*, 70(1):191–221.

Barakchian, S. M. and Crowe, C. (2013). Monetary policy matters: Evidence from new shocks data. *Journal of Monetary Economics*, 60(8):950 – 966.

Barigozzi, M., Conti, A. M., and Luciani, M. (2014). Do euro area countries respond asymmetrically to the common monetary policy? *Oxford Bulletin of Economics and Statistics*, 76(5):693–714.

Berben, R.-P., Locarno, A., Morgan, J., and Valls, J. (2004). Cross-country differences in monetary policy transmission. Working Paper Series 400, European Central Bank.

Bernanke, B. S., Boivin, J., and Eliasz, P. (2005). Measuring the effects of monetary policy: A factor-augmented vector autoregressive (FAVAR) approach. *The Quarterly Journal of Economics*, 120(1):387–422.

Boivin, J. and Giannoni, M. (2007). DSGE models in a data-rich environment. Working papers 162, Banque de France.

Calza, A., Monacelli, T., and Stracca, L. (2013). Housing Finance And Monetary Policy. *Journal of the European Economic Association*, 11:101–122.

Campbell, J., Evans, C., Fisher, J., and Justiniano, A. (2012). Macroeconomic effects of Federal Reserve forward guidance. 2012.

Ciccarelli, M., Maddaloni, A., and Peydró, J.-L. (2013). Heterogeneous transmission mechanism: monetary policy and financial fragility in the eurozone. *Economic Policy*, 28(75):459–512.

Cochrane, J. H. and Piazzesi, M. (2002). The Fed and interest rates - a high-frequency identification. *American Economic Review*, 92(2):90–95.

Duarte, J. B. and Dias, D. A. (2016). Housing and monetary policy in the business cycle: What do housing rents have to say? *mimeo*, 11.

ECB (2017). Financial integration in Europe. Technical report.

Faust, J., Swanson, E. T., and Wright, J. H. (2004). Identifying VARS based on high frequency futures data. *Journal of Monetary Economics*, 51(6):1107 – 1131.

Forni, M. and Gambetti, L. (2010). The dynamic effects of monetary policy: A structural factor model approach. *Journal of Monetary Economics*, 57(2):203 – 216.

Gertler, M. and Karadi, P. (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, 7(1):44–76.

Giannone, D., Reichlin, L., and Sala, L. (2005). Monetary policy in real time. In *NBER Macroeconomics Annual 2004*, volume 19, pages 161–224. MIT Press.

Gurkaynak, R. S., Sack, B., and Swanson, E. T. (2005). Do actions speak louder than words? the response of asset prices to monetary policy actions and statements. *International Journal of Central Banking*, 1(1):55–93.

Hamilton, J. D. (2008). Daily monetary policy shocks and new home sales. *Journal of Monetary Economics*, 55(7):1171 – 1190.

Hanson, S. G. and Stein, J. C. (2015). Monetary policy and long-term real rates. *Journal of Financial Economics*, 115(3):429 – 448.

Iacoviello, M. (2005). House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle. *American Economic Review*, 95(3):739–764.

Jarocinski, M. and Karadi, P. (2018). Deconstructing monetary policy surprises: the role of information shocks. *ECB Working Paper Series*, (2133).

Kaplan, G., Moll, B., and Violante, G. L. (2016). Monetary policy according to HANK. *National Bureau of Economic Research*.

Kim, C. H. and Other, L. (2017). The short-run effect of monetary policy shocks on credit risk - an analysis of the Euro Area.

Kuttner, K. N. (2001). Monetary policy surprises and interest rates: Evidence from the Fed funds futures market. *Journal of Monetary Economics*, 47(3):523 – 544.

Lloyd, S. (2017a). Estimating nominal interest rate expectations: Overnight indexed swaps and the term structure. Cambridge working papers in economics, University of Cambridge.

Lloyd, S. (2017b). Overnight indexed swap market-based measures of monetary policy expectations. Cambridge working papers in economics, University of Cambridge.

Mertens, K. and Ravn, M. O. (2013). The dynamic effects of personal and corporate income tax changes in the United States. *American Economic Review*, 103(4):1212–47.

Mian, A., Rao, K., and Sufi, A. (2013). Household balance sheets, consumption, and the economic slump. *Quarterly Journal of Economics*, 128(4).

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

Nakamura, E. and Steinsson, J. (2013). High frequency identification of monetary non-neutrality: The information effect. Working Paper 19260, National Bureau of Economic Research.

Onatski, A. (2009). Testing hypotheses about the number of factors in large factor models. *Econometrica*, 77(5):1447–1479.

Onatski, A. (2010). Determining the number of factors from empirical distribution of eigenvalues. *The Review of Economics and Statistics*, 92(4):1004–1016.

Osborne, J. (2005). Housing in the euro area - twelve markets one money. Quarterly Bulletin 4, Central Bank of Ireland.

Owen, A. B. and Wang, J. (2015). Bi-cross-validation for factor analysis. *ArXiv e-prints*.

Romer, C. D. and Romer, D. H. (2002). A rehabilitation of monetary policy in the 1950's. *American Economic Review*, 92(2):121–127.

Sargent, T. and Sims, C. (1977). Business cycle modeling without pretending to have too much a priori economic theory. Working Papers 55, Federal Reserve Bank of Minneapolis.

Sargent, T. J. (1989). Two models of measurements and the investment accelerator. *Journal of Political Economy*, 97(2):251–287.

Stock, J. H. and Watson, M. W. (2005). Implications of dynamic factor models for VAR analysis. Working Paper 11467, National Bureau of Economic Research.

Stock, J. H. and Watson, M. W. (2012). Disentangling the channels of the 2007-2009 recession. NBER Working Papers 18094, National Bureau of Economic Research.

Stock, J. H. and Watson, M. W. (2016). Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics. volume 2 of *Handbook of Macroeconomics*, pages 415 – 525. Elsevier.

Westig, D. and Bertalot, L. (2016). Hypostat 2016: A review of Europe's housing and mortgage markets. Technical report, European Mortgage Federation.

A Selecting the Number of Factors - Additional Figures
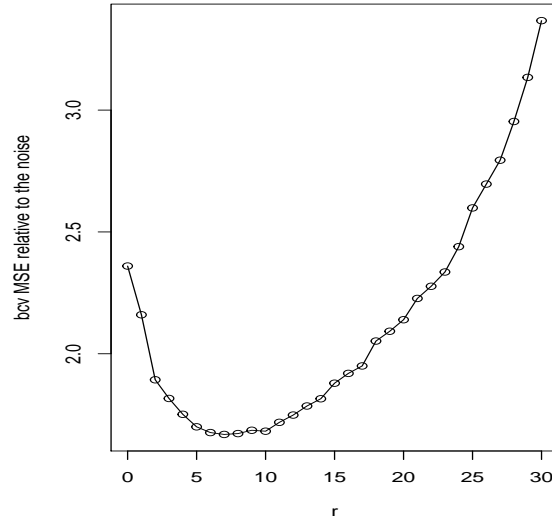
Figure 12: Bi-cross-validation method proposed by Owen and Wang (2015)
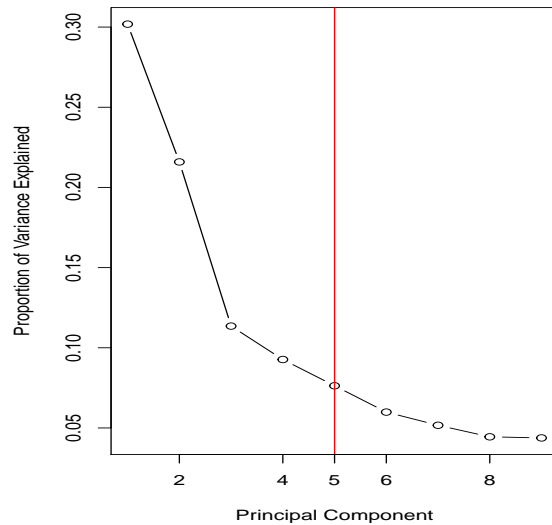
Figure 13: Variance explained by each additional factor

B Small VAR with High-Frequency Identification

In this section we use our instrument to identify monetary policy shocks in a simple VAR with three variables: output, consumer prices and a policy indicator. This simpler setting is useful to test the strength of the external instrument. Estimating a simple VAR for monthly and quarterly data, we test different instruments and policy indicators. The set of instruments to be tested comprises 3-month, 6-month and 12-month EONIA futures. The set of policy indicators

Figure 14: VAR using monthly data from 2000 to 2016. Here we show the responses to a one standard deviation shock in the policy indicator, comparing the high-frequency identification with a Cholesky identification strategy. The dashed lines report the bootstrapped 68% confidence intervals. The Cholesky identification orders the policy indicator last. The F-test for the first-stage regression on the external instrument is 4.85 and the R^2 is 2 percent.

is given by EONIA, one-year aggregate EA bond yields, one-year German government bond yields, as well as two-year German government bond yields. We use industrial production (IP) as a measure of output for monthly data, and real GDP for quarterly data. For consumer prices, we use HICP at both frequencies.

The combination of policy indicator and instrument that provides the best instrument strength is the one selected to report the dynamic effects of monetary policy shocks on output and consumer prices. For monthly data, the selected instrument is the 3-month EONIA future and the policy indicator is the two-year German government bond rate, while for the quarterly data the instrument that works best is the one-year EONIA future and the policy indicator is the one-year German government bond rate.

In order to compare our identification strategy for the EA with a more standard identification, we also estimate the impulse-response functions using the Cholesky decomposition with the following ordering: output, consumer prices and policy indicator. The results with monthly data are reported in Figure 14. The more traditional approach to identify monetary policy surprises

exhibits both a price puzzle and an output puzzle. Interestingly, when using our external instrument approach, both puzzles disappear. The external instrument delivers responses that are more in line with standard economic theory where output falls temporarily and recovers in the medium-run (neutrality), and prices fall. In this specification, the instrument is weak as its F-test is below 10 which implies the possibility of biased estimates in a small sample such as ours. However, in the case of a just identified IV, it is possible to get approximately unbiased (or less biased) estimates even with weak instruments.

Using quarterly data, we get a significantly stronger instrument with a first-stage F-test of 19.45. Figure 15 shows the same set of variable responses, now using quarterly data. The Cholseky identification does not feature a price puzzle in this setup. There is, however, an output puzzle. With the high-frequency identification, on the other hand, we only get a price puzzle on the contemporaneous response, while there is no output puzzle. The limitations of an identification strategy based on timing restrictions are further highlighted at the quarterly frequency as it is hard to argue that consumer prices (collected on a monthly basis) do not react in the same quarter to monetary policy surprises. If we want to allow prices to respond contemporaneously, we can order consumer prices last (instead of the monetary policy indicator). However, in this case we also get the undesirable restriction of not letting monetary policy react to consumer prices contemporaneously. The external instrument is able circumvent this limitation.

Figure 16 shows the responses when we order the consumer prices last in the Cholesky decomposition. In this case, consumer prices are allowed to react contemporaneously to monetary policy shocks. When the consumer price response is not contemporaneously restricted to zero, we find that the price puzzle is present and, contrary to the high-frequency identification, it lasts for a few quarters after the shock hits the economy.
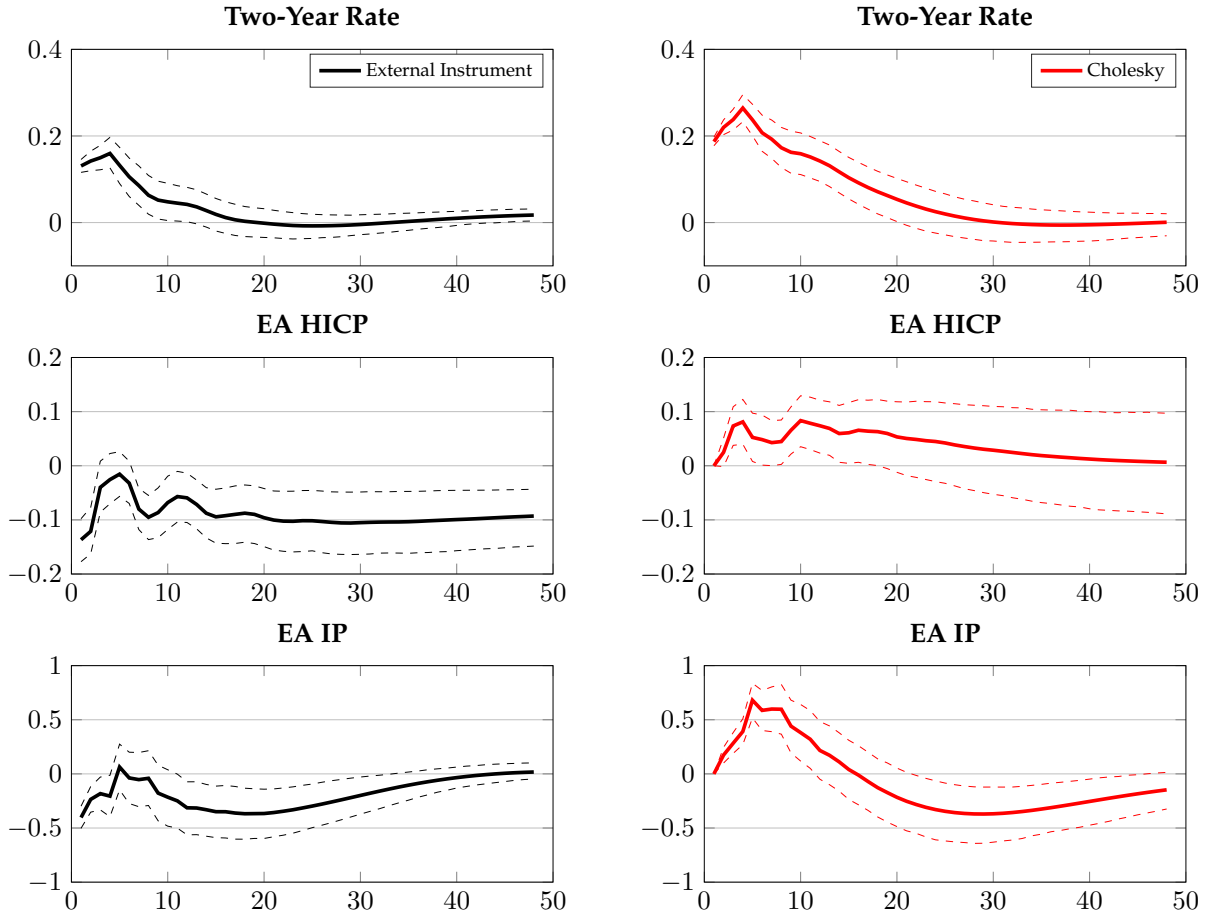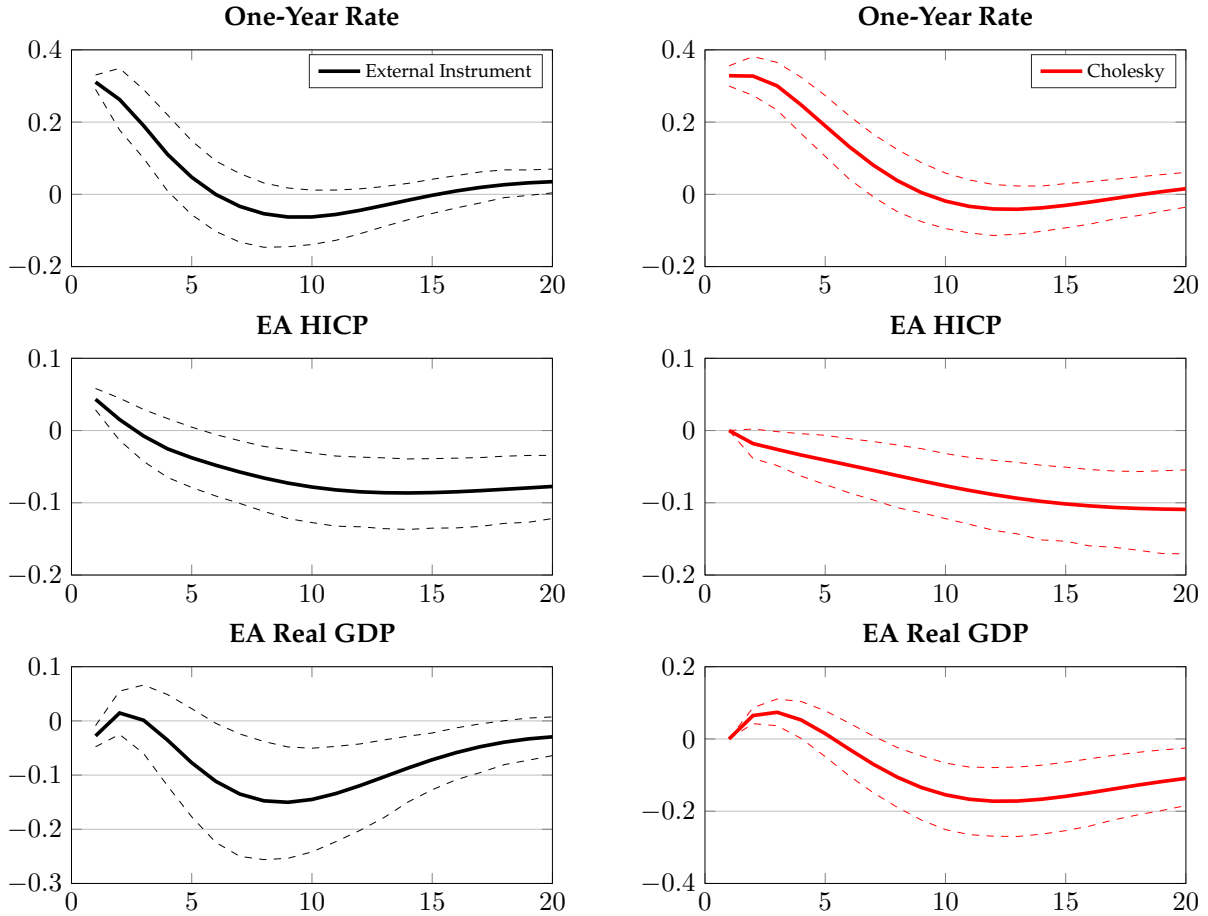
Figure 15: VAR using quarterly data from 2000 to 2016. Here we show the responses to a one standard deviation shock in the policy indicator using the high-frequency identification and the Cholesky identification. The dashed lines report the bootstrapped 68% confidence intervals. The Cholesky identification orders the policy indicator last. The F-test for the first-stage regression on the external instrument is 19.45 and the R^2 is 22 percent.
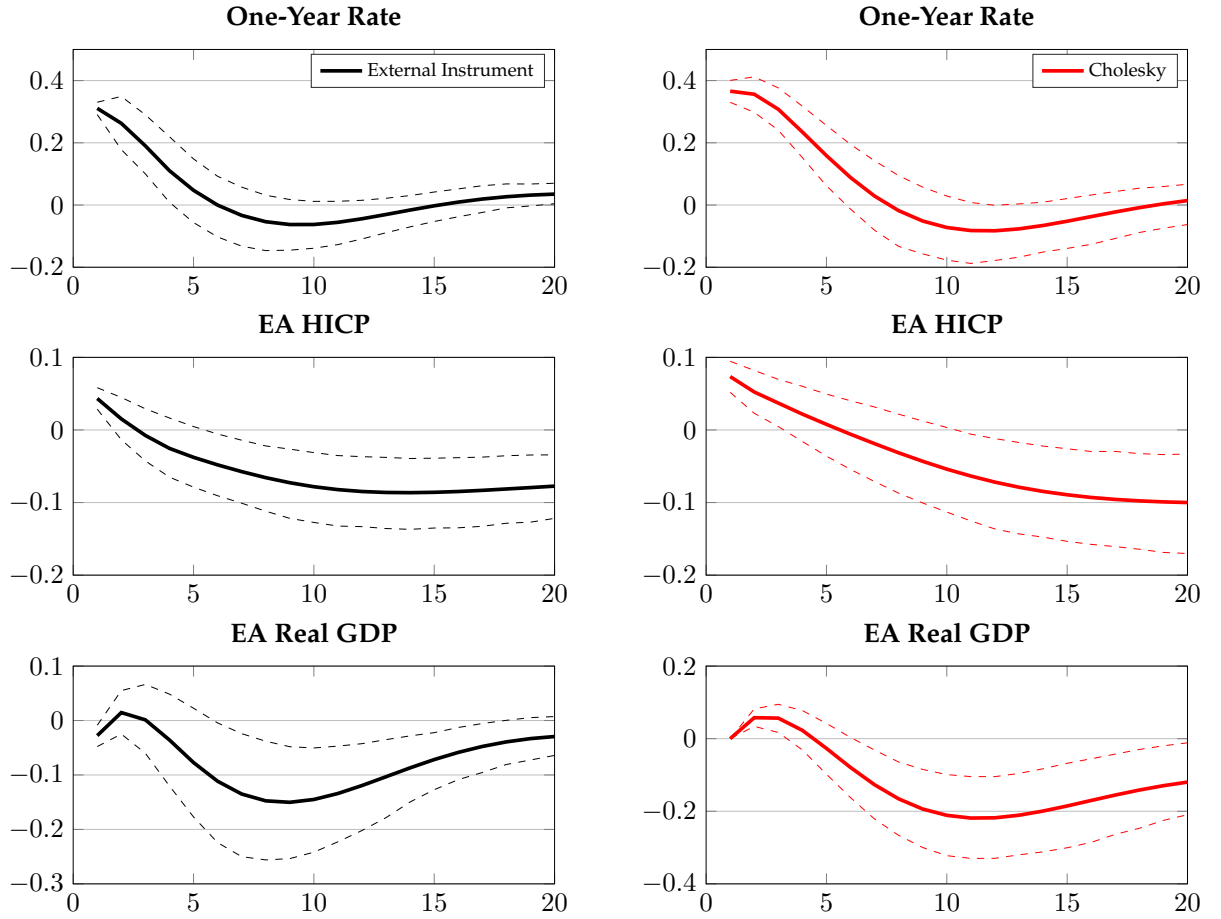
Figure 16: VAR using quarterly data from 2000 to 2016. Here we show the responses to a one standard deviation shock in the policy indicator using high-frequency and Cholesky identification. The dashed lines report the bootstrapped 68% confidence intervals. Here, the Cholesky identification orders the consumer prices last. The F-test for the first-stage regression on the external instrument is 19.45 and the R^2 is 22 percent.

C Data Set

Table 5 contains a complete list of the series in our data set as well as detailed descriptions and information regarding transformations, geographical coverage and sources. Abbreviations and codes are laid out in the following:

Transformation code (T)

1 - no transformation

2 - difference in levels

4 - logs

5 - difference in logs

Geography

EA - Euro area

EA12 - Euro area (12 countries)

EA19 - Euro area (19 countries)

EACC - Euro area (changing composition)

EA11_i - 11 individual series for sample countries

Factor analysis (F)

Y - included in data set for principal component analysis

Seasonal adjustment

WDSA - working day and seasonally adjusted

SA - seasonally adjusted

NA - neither working day nor seasonally adjusted

Note: National house price indices have different start dates across countries. They begin in 2005 Q4 for Spain, 2006 Q2 for France, 2007 Q1 for Luxembourg, 2008 Q1 for Portugal, 2010 Q1 for Italy and Austria, and 2005 Q1 for all other countries. Furthermore, unemployment data for France between 2000 Q1 and 2005 Q1, as well as Luxembourg between 2000 Q1 and 2003 Q1 is only available annually and has been linearly interpolated to create a quarterly data series. Thereafter all unemployment data is quarterly. Finally, import and export data for Germany, Spain and Italy is only available from 2012 Q1 onward.

	Description	T	Source	Geography	Start	End	F
GDP & Personal Income							
GDP	Gross Domestic Product at market prices, Chain linked volumes, 2010=100, WDSA	5	Eurostat	EA12	2000 Q1	2016 Q4	
PCON	Household and NPISH final consumption expenditure, Chain linked volumes, 2010=100, WDSA	5	Eurostat	EA12	2000 Q1	2016 Q4	
G	Final consumption expenditure of general government, Chain linked volumes, 2010=100, WDSA	5	Eurostat	EA12	2000 Q1	2016 Q4	
GFCF	Gross fixed capital formation, Chain linked volumes, 2010=100, WDSA	5	Eurostat	EA12	2000 Q1	2016 Q4	
EX	Exports of goods and services, Chain linked volumes, 2010=100, WDSA	5	Eurostat	EA12	2000 Q1	2016 Q4	
IM	Imports of goods and services, Chain linked volumes, 2010=100, WDSA	5	Eurostat	EA12	2000 Q1	2016 Q4	
GDP.i	Gross Domestic Product at market prices, Chain linked volumes, 2010=100, WDSA	5	Eurostat	EA11.i	2000 Q1	2016 Q4	
CON.i	Final consumption expenditure, Chain linked volumes, 2010=100, WDSA	5	Eurostat	EA11.i	2000 Q1	2016 Q4	
PCON.i	Household and NPISH final consumption expenditure, Chain linked volumes, 2010=100, WDSA	5	Eurostat	EA11.i	2000 Q1	2016 Q4	Y
G.i	Final consumption expenditure of general government, Chain linked volumes, 2010=100, WDSA	5	Eurostat	EA11.i	2000 Q1	2016 Q4	Y
GFCF.i	Gross fixed capital formation, Chain linked volumes (2010), million euro, WDSA	5	Eurostat	EA11.i	2000 Q1	2016 Q4	Y
EX.i	Exports of goods and services, Chain linked volumes, 2010=100, unadjusted data	5	Eurostat	EA11.i	2000 Q1	2016 Q4	Y
IM.i	Imports of goods and services, Chain linked volumes, 2010=100, unadjusted data	5	Eurostat	EA11.i	2000 Q1	2016 Q4	Y
Prices/Deflators							
GDPDEF	Gross domestic product at market prices, Price index (implicit deflator), 2010=100, euro, WDSA	5	Eurostat	EA12	2000 Q1	2016 Q4	
PCONDEF	Household and NPISH final consumption expenditure, Price index (implicit deflator), 2010=100, euro, WDSA	5	Eurostat	EA12	2000 Q1	2016 Q4	
GDEF	Final consumption expenditure of general government, Price index (implicit deflator), 2010=100, euro, WDSA	5	Eurostat	EA12	2000 Q1	2016 Q4	
GFCFDEF	Gross fixed capital formation, Price index (implicit deflator), 2010=100, euro, WDSA	5	Eurostat	EA12	2000 Q1	2016 Q4	
EXDEF	Exports of goods and services, Price index (implicit deflator), 2010=100, euro, WDSA	5	Eurostat	EA12	2000 Q1	2016 Q4	Y
IMDEF	Imports of goods and services, Price index (implicit deflator), 2010=100, euro, WDSA	5	Eurostat	EA12	2000 Q1	2016 Q4	Y
PPI	Producer prices in industry, domestic market, index 2010=100, unadjusted data	5	Eurostat	EA19	2000 Q1	2016 Q4	
HICP00	All-items HICP, Index, 2015=100	5	Eurostat	EACC	2000 Q1	2016 Q4	
HICP01	HICP Food and non-alcoholic beverages, Index, 2015=100	5	Eurostat	EACC	2000 Q1	2016 Q4	
HICP02	HICP Alcoholic beverages, tobacco and narcotics, Index, 2015=100	5	Eurostat	EACC	2000 Q1	2016 Q4	
HICP03	HICP Clothing and footwear, Index, 2015=100	5	Eurostat	EACC	2000 Q1	2016 Q4	
HICP05	HICP Furnishings, household equipment and routine household maintenance, Index, 2015=100	5	Eurostat	EACC	2000 Q1	2016 Q4	

Code	Description	Y		Source	Region	Start	End
HICP06	HICP Health, Index, 2015=100		5	Eurostat	EACC	2000 Q1	2016 Q4
HICP07	HICP Transport, Index, 2015=100		5	Eurostat	EACC	2000 Q1	2016 Q4
HICP08	HICP Communications, Index, 2015=100		5	Eurostat	EACC	2000 Q1	2016 Q4
HICP09	HICP Recreation and culture, Index, 2015=100		5	Eurostat	EACC	2000 Q1	2016 Q4
HICP10	HICP Education, Index, 2015=100		5	Eurostat	EACC	2000 Q1	2016 Q4
HICP11	HICP Restaurants and hotels, Index, 2015=100		5	Eurostat	EACC	2000 Q1	2016 Q4
HICP12	HICP Miscellaneous goods and services, Index, 2015=100		5	Eurostat	EACC	2000 Q1	2016 Q4
HICPXFD	Overall HICP index excluding seasonal food, Index, 2015=100		5	Eurostat	EACC	2000 Q4	2016 Q4
HICPXUTIL	Overall HICP index excluding housing, water, electricity, gas and other fuels, 2015=100		5	Eurostat	EACC	2000 Q1	2016 Q4
HICPXHTH	Overall HICP index excluding education, health and social protection, Index, 2015=100		5	Eurostat	EACC	2000 Q1	2016 Q4
HICPUTIL	HICP Housing, water electricity, gas and other fuels, Index, 2015=100		5	Eurostat	EACC	2000 Q1	2016 Q4
PPIING	Producer Price Index, MIG - intermediate goods, unadjusted data, 2010=100		5	Eurostat	EA19	2000 Q1	2016 Q4
PPICAG	Producer Price Index, MIG - capital goods, unadjusted data, 2010=100		5	Eurostat	EA19	2000 Q1	2016 Q4
PPINDCOG	Producer Price Index, MIG - non-durable consumer goods, unadjusted data, 2010=100		5	Eurostat	EA19	2000 Q1	2016 Q4
PPIM	Producer Price Index, Manufacturing, unadjusted data, 2010=100		5	Eurostat	EA19	2000 Q1	2016 Q4
HICP_i	Individual country HICP	Y	5	Eurostat	EA11_i	2000 Q1	2016 Q4
UTIL_i	HICP Housing, water electricity, gas and other fuels, Index, 2015=100		5	Eurostat	EA11_i	2000 Q1	2016 Q4
PPI_i	Producer prices in industry (except construction sewerage, waste management and remediation activities), Domestic output price index in national currency, 2010=100	Y	5	Eurostat	EA11_i	2000 Q1	2016 Q4
CDEF_i	Final consumption expenditure, Price index (implicit deflator), 2010=100, euro, WDSA	Y	5	Eurostat	EA11_i	2000 Q1	2016 Q4
PCONDEF_i	Household and NPISH final consumption expenditure, Price index (implicit deflator), 2010=100, euro, WDSA		5	Eurostat	EA11_i	2000 Q1	2016 Q4
GFCFDEF_i	Gross fixed capital formation, Price index (implicit deflator), 2010=100, euro, WDSA	Y	5	Eurostat	EA11_i	2000 Q1	2016 Q4
CPIIMF	IMF World Commodity Price Index, USD denominated, weights based on 2002-2004 average world export earnings, non-fuel primary commodities and energy, 2005=100		5	IMF	World	2000 Q1	2016 Q4
CPIECB	ECB Commodity Price Index, Euro denominated, use-weighted, Total non-energy commodity, unadjusted data, 2010=100	Y	5	ECB SDW	EA19	2000 Q1	2016 Q4
OIL	Brent crude oil 1-month forward, fob (free on board) per barrel, Euro	Y	5	ECB SDW	EACC	2000 Q1	2016 Q4

Industrial Production

Code	Description	Y		Source	Region	Start	End
IPIT	Industrial Production Index, Total Industry, WDSA, 2005=100		5	ECB SDW	EA19	2000 Q1	2016 Q4
IPIING	Industrial Production Index, MIG - intermediated goods, WDSA, 2010=100		5	Eurostat	EA19	2000 Q1	2016 Q4

42

Code	Description		Source	Region	Start	End	
IPINRG	Industrial Production Index, MIG - energy, WDSA, 2010=100	5	Eurostat	EA19	2000 Q1	2016 Q4	Y
IPICAG	Industrial Production Index, MIG - capital goods, WDSA, 2010=100	5	Eurostat	EA19	2000 Q1	2016 Q4	Y
IPICOG	Industrial Production Index, MIG - consumer goods, WDSA, 2010=100	5	Eurostat	EA19	2000 Q1	2016 Q4	
IPIDCOG	Industrial Production Index, MIG - durable consumer goods, WDSA, 2010=100	5	Eurostat	EA19	2000 Q1	2016 Q4	Y
IPINDCOG	Industrial Production Index, MIG - non-durable consumer goods, WDSA, 2010=100	5	Eurostat	EA19	2000 Q1	2016 Q4	Y
IPIMQ	Industrial Production Index, Mining and quarrying, WDSA, 2010=100	5	Eurostat	EA19	2000 Q1	2016 Q4	
IPIM	Industrial Production Index, Manufacturing, WDSA, 2010=100	5	Eurostat	EA19	2000 Q1	2016 Q4	
ITIING	Industrial Turnover Index, MIG Intermediate Goods (2010=100, WDSA)	5	Eurostat	EA19	2000 Q1	2016 Q4	Y
ITINRG	Industrial Turnover Index, MIG Energy (2010=100, WDSA)	5	Eurostat	EA19	2000 Q1	2016 Q4	Y
ITICAG	Industrial Turnover Index, MIG Capital Goods (2010=100, WDSA)	5	Eurostat	EA19	2000 Q1	2016 Q4	Y
ITICOG	Industrial Turnover Index, MIG Consumer Goods (2010=100, WDSA)	5	Eurostat	EA19	2000 Q1	2016 Q4	
ITIDCOG	Industrial Turnover Index, MIG Durable Consumer Goods (2010=100, WDSA)	5	Eurostat	EA19	2000 Q1	2016 Q4	Y
ITINDCOG	Industrial Turnover Index, MIG Non-Durable Consumer Goods (2010=100, WDSA)	5	Eurostat	EA19	2000 Q1	2016 Q4	Y
CAPUTIL	Current level of capacity utilization, percent	1	Eurostat	EA19	2000 Q1	2016 Q4	Y
ITIM	Industrial Turnover Index, Manufacturing, 2010=100, SWDA	5	Eurostat	EA19	2000 Q1	2016 Q4	

Employment and Unemployment

Code	Description		Source	Region	Start	End	
WIN	Compensation of employees, Current prices, million euro, WDSA	5	Eurostat	EA12	2000 Q1	2016 Q4	Y
U	Total Unemployment rate (quarterly average), WDSA	2	Eurostat	EACC	2000 Q1	2016 Q4	
EMP	Total employment	5	Eurostat	EA19	2000 Q1	2016 Q4	Y
U_i	Unemployment rate, total from age 15 to 74, percentage	2	Eurostat	EA11.i	2000 Q1	2016 Q4	Y
LABCON	Labour Input in Construction, Index of Hours Worked, 2010=100, WDSA	5	Eurostat	EA19	2000 Q1	2016 Q4	Y

Housing Starts

Code	Description		Source	Region	Start	End	
BUILD	Building Permits, Residential Buildings, Index, 2010=100, WDSA	5	Eurostat	EA19	2000 Q1	2016 Q4	Y
GFCFC	Gross fixed capital formation: Total construction (gross), chain linked volumes, Index, 2010=100	5	Eurostat	EACC	2000 Q1	2016 Q4	
GFCFC.i	Gross fixed capital formation: Total construction (gross), chain linked volumes, Index, 2010=100	5	Eurostat	11 ex BEL	2000 Q1	2016 Q4	
GFCFD	Gross fixed capital formation: Dwellings (gross), chain linked volumes, Index, 2010=100	5	Eurostat	EACC	2000 Q1	2016 Q4	
GFCFD.i	Gross fixed capital formation: Dwellings (gross), chain linked volumes, Index, 2010=100	5	Eurostat	11 ex BEL	2000 Q1	2016 Q4	
PROCO	Production in Construction, Volume Index, 2010=100, WDSA	5	Eurostat	EA19	2000 Q1	2016 Q4	Y

Inventories, Orders and Sales

Code	Description		Source	Region	Start	End	
ORDING	Industrial New Orders, MIG Intermediate Goods, 2010=100, SA	5	ECB SDW	EA19	2000 Q1	2016 Q4	
ORDCAG	Industrial New Orders, MIG Capital Goods, 2010=100, SA	5	ECB SDW	EA19	2000 Q1	2016 Q4	Y

Code	Description		Source	Area			
ORDCOG	Industrial New Orders, MIG Consumer Goods, 2010=100, SA	5	ECB SDW	EA19	2000 Q1	2016 Q4	Y
ORDM	Industrial New Orders, Manufacturing, 2010=100, SA	5	ECB SDW	EA19	2000 Q1	2016 Q4	
Earnings and Productivity							
PRD.i	Real labour productivity per hour worked, 2010=100, unadjusted data	5	Eurostat	EA11.i	2000 Q1	2016 Q4	Y
ULC.i	Nominal unit labour cost based on hours worked, 2010=100, unadjusted data	5	Eurostat	11 ex BEL	2000 Q1	2016 Q4	
Money and Credit							
EUSWE1	1 year EONIA swap	1	Bloomberg	EA	2000 Q1	2016 Q4	
STINT	3-month money market interest rate	1	Eurostat	EACC	2000 Q1	2016 Q4	
LTINT	EMU convergence criterion long-term bond yields	1	Eurostat	EACC	2000 Q1	2016 Q4	
MIR	Bank interest rates - loans to households for house purchase (outstanding amount business coverage), average of observations through period, percent per annum	1	ECB SDW	EACC	2003 Q1	2016 Q4	
COB	Cost of borrowing for households for house purchase (new business coverage), average of observations through period, percent per annum	1	ECB SDW	EACC	2003 Q1	2016 Q4	
EURIBOR3MD	3-Month Euro Interbank Offered Rate (%, NSA)	1	ECB SDW	EA	2000 Q1	2016 Q4	
EURIBOR6MD	6-Month Euro Interbank Offered Rate (%, NSA)	1	ECB SDW	EA	2000 Q1	2016 Q4	
EURIBOR1YD	1-Year Euro Interbank Offered Rate (%, NSA)	1	ECB SDW	EA	2000 Q1	2016 Q4	
YLD.3Y	3-Year Euro Area Government Benchmark Bond Yield (%, NSA)	1	ECB SDW	EA	2004 Q4	2016 Q4	
YLD.5Y	5-Year Euro Area Government Benchmark Bond Yield (%, NSA)	1	ECB SDW	EA	2004 Q4	2016 Q4	
YLD.10Y	10-Year Euro Area Government Benchmark Bond Yield (%, NSA)	1	ECB SDW	EA	2004 Q4	2016 Q4	
EONIA	Euro Overnight Index Average (%, NSA)	1	ECB SDW	EA	2000 Q1	2016 Q4	Y
REFI	ECB Official Refinancing Operation Rate (effective, %, NSA)	1	ECB SDW	EA	2000 Q1	2016 Q4	
S3MDREFI	Spread EURIBOR3MD - REFI	1	ECB SDW	EA	2000 Q1	2016 Q4	
S10YYLDREFI	Spread YLD_10Y - REFI	1	ECB SDW	EA	2004 Q4	2016 Q4	
LTINT.i	Long-term interest rates, percent per annum	1	OECD	EA11.i	2000 Q1	2016 Q4	Y
STINT.i	Short-term interest rates, percent per annum	1	OECD	EA11.i	2000 Q1	2016 Q4	Y
MIR.i	Bank interest rates - loans to households for house purchase (outstanding amount business coverage), average of observations through period, percent per annum	1	ECB SDW	EA11.i	2003 Q1	2016 Q4	
COB.i	Cost of borrowing for households for house purchase (new business coverage), average of observations through period, percent per annum	1	ECB SDW	EA11.i	2003 Q1	2016 Q4	
Stock Prices, Wealth, Household Balance Sheets							
SP	Share prices, Index, 2010=100	5	OECD	EA19	2000 Q1	2016 Q4	

Code	Description		Source	Region	Start	End	
SP_i	Share prices, Index, 2010=100	5	OECD	EA11.i	2000 Q1	2016 Q4	Y
OWN_i	Distribution of population by tenure status: ownership, percentage	5	Eurostat	EA11.i	2003 Q1	2016 Q4	

Housing Prices

Code	Description		Source	Region	Start	End	
RENTS	HICP Actual rentals for housing, Index, 2015=100	5	Eurostat	EACC	2000 Q1	2016 Q4	
HPI	House price index, 2010=100	5	Eurostat	EACC	2005 Q1	2016 Q4	
RHPI	Real house prices (=HPI/HICP00)	5	Eurostat	EACC	2005 Q1	2016 Q4	
RRENTS	Real rents (=RENTS/HICP00)	5	Eurostat	EACC	2000 Q1	2016 Q4	
BUILDCOSTI	Construction Cost Index, Residential Buildings (2010=100, WDSA)	5	Eurostat	EA19	2000 Q1	2016 Q4	
REN_i	HICP Actual rentals for housing, Index, 2015=100	5	Eurostat	EA11.i	2000 Q1	2016 Q4	
RREN_i	Real rents (=REN/HICP00)	5	Author's calculation	EA11.i	2000 Q1	2016 Q4	
HPL_i	House price index, 2010=100	5	Eurostat	EA11.i	2005 Q1	2016 Q4	
RHPL_i	Real house prices (=HPI/HICP00)	5	Author's calculation	EA11.i	2005 Q1	2016 Q4	
NDW_i	House price index, New dwellings, 2010=100	5	Eurostat	11 ex NLD	2005 Q1	2016 Q4	
EDW_i	House price index, Existing dwellings, 2010=100	5	Eurostat	EA11.i	2005 Q1	2016 Q4	

Exchange Rates

Code	Description		Source	Region	Start	End	
NEER	Euro Nominal Effective Exchange Rate - 42 trading partners, Index, 2005=100	5	Eurostat	EA19	2000 Q1	2016 Q4	Y
EXRUK	Foreign Exchange Rate: United Kingdom (GBP per EUR - quarterly average)	5	Eurostat	EA	2000 Q1	2016 Q4	Y
EXRSW	Foreign Exchange Rate: Switzerland (CHF per EUR - quarterly average)	5	Eurostat	EA	2000 Q1	2016 Q4	Y
EXRJP	Foreign Exchange Rate: Japan (JPY per EUR - quarterly average)	5	Eurostat	EA	2000 Q1	2016 Q4	Y
EXRUS	Foreign Exchange Rate: United States of America (USD per EUR - quarterly average)	5	Eurostat	EA	2000 Q1	2016 Q4	Y

Expectations

Code	Description		Source	Region	Start	End	
BSBCI	EA Business Climate Indicator (SA)	2	Eurostat	EA19	2000 Q1	2016 Q4	
BSCCI	Construction Confidence Indicator (SA)	2	Eurostat	EA19	2000 Q1	2016 Q4	
BSESI	Economic Sentiment Indicator (SA)	2	Eurostat	EA19	2000 Q1	2016 Q4	
BSICI	Industrial Confidence Indicator (SA)	2	Eurostat	EA19	2000 Q1	2016 Q4	
BSRCI	Retail Confidence Indicator (SA)	2	Eurostat	EA19	2000 Q1	2016 Q4	
BSCSMCI	Consumer Confidence Indicator (SA)	2	Eurostat	EA19	2000 Q1	2016 Q4	Y
BSSCI	Services Confidence Indicator (SA)	2	Eurostat	EA19	2000 Q1	2016 Q4	

D On Interpreting Factors

For Table 6, we regress each transformed data series on one of the 5 factors at a time and subsequently report the series where these regression resulted in the highest R^2. While by nature principal component analysis does not identify factors economically, the table gives a rough indication of the information represented by them. On this basis, we suggest the following tentative interpretation:

Factor 1 is likely to represent prices in the economy. It shows a high correlation with a variety of price indices, from producer prices to HICP, and explains over half of the variance in these series. Factor 2 is very closely related to measures of interest rates. This includes money-market rates, as well as borrowing rates for house purchase. Factors 3 and 4 appear to contain a substantial amount of information about labour markets, with high correlations to unit labour cost and unemployment rates. That said, the factors are also closely related to other variables and an interpretation seem much more contentious than for factors 1 and 2. Factor 5 picks up information from various areas of macroeconomic activity and we do not believe that a straightforward interpretation of the factor is possible.

On the whole, we can emphasise that factors 1 and 2 seem to represent the economic concepts of *prices* and *interest rates*. More generally, the latter could also be interpreted as representing *financial conditions*.

Table 6: List of series that are best explained by a single extracted factor according to R-squared of a linear regression of the (transformed) series on the respective factor.

	Series	R-squared
	Producer Prices in Industry	0.67
	Harmonised Index of Consumer Prices	0.56
Factor 1	Industrial Turnover Index, Manufacturing	0.53
	Compensation of Employees	0.49
	Gross Fixed Capital Formation Price Index	0.48
	Cost of Borrowing for Households for House Purchase	0.49
	6-month Euribor	0.45
Factor 2	1-year Euribor	0.45
	3-month Euribor	0.44
	Long-term Interest Rate Belgium	0.43
	Government Spending Italy	0.61
	Unit Labour Cost Germany	0.61
Factor 3	Government Spending Finland	0.61
	Unit Labour Cost Luxembourg	0.60
	Unit Labour Cost Italy	0.60
	Unemployment Italy	0.63
	Unemployment Netherlands	0.49
Factor 4	Real House Prices Ireland	0.44
	Unemployment Finland	0.43
	Real House Prices France	0.43
	Real House Prices Netherlands	0.46
	GDP Spain	0.40
Factor 5	Private Consumption Spain	0.33
	House Prices Netherlands	0.32
	Gross Fixed Capital Formation in Construction	0.32

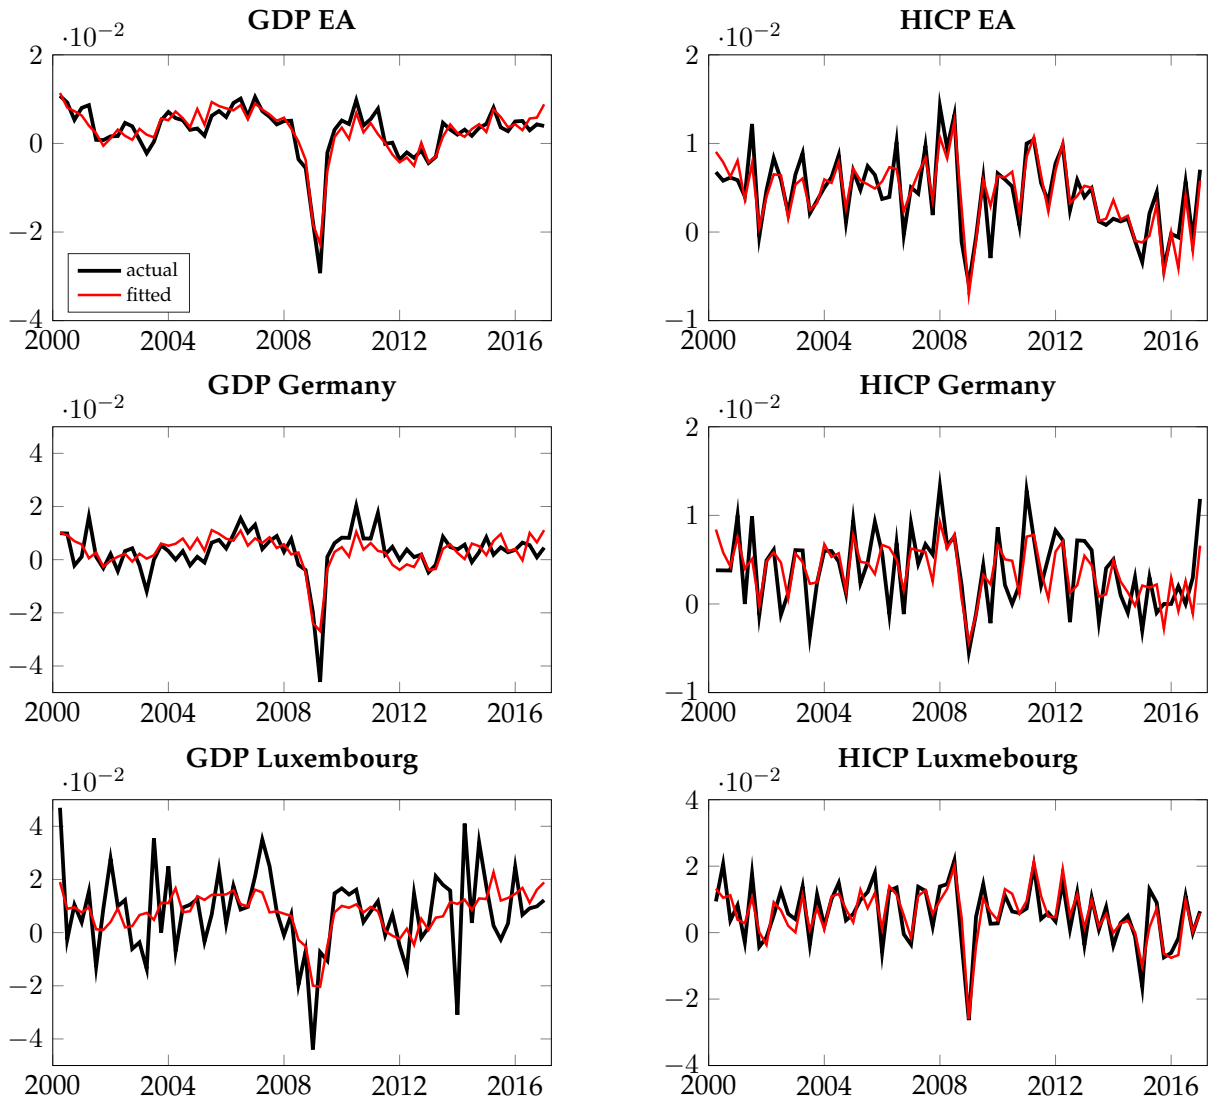# E Explanatory Power of Factors

Figure 17: The figure compares actual (transformed) GDP and HICP data with corresponding fitted series on the basis of 5 extracted factors for the euro area (EA), Germany and Luxembourg from 2000 Q1 to 2016 Q4. Germany and Luxembourg represent the largest and smallest economies in our sample, respectively. In DFM terminology, the fitted series represent the *systematic* component of the data series, while the actual series also contains an *idiosyncratic* component.

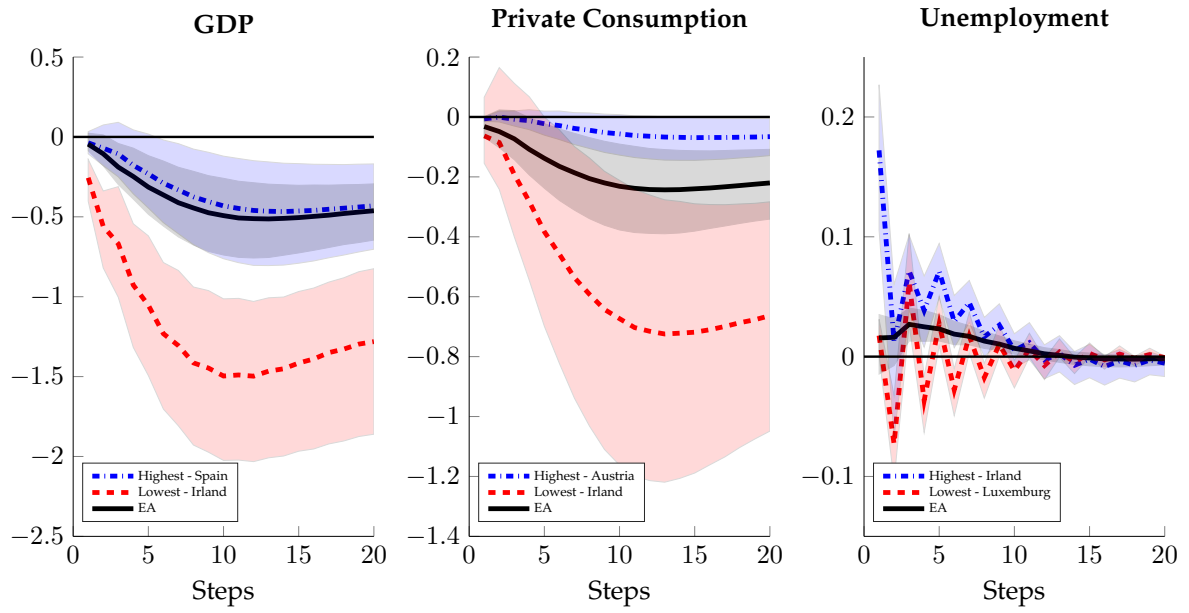# F Highest and lowest responses to monetary policy shock

Figure 18: Highest/lowest percentage responses of selected real variables to a 25bp contractionary policy shock across euro area member countries.
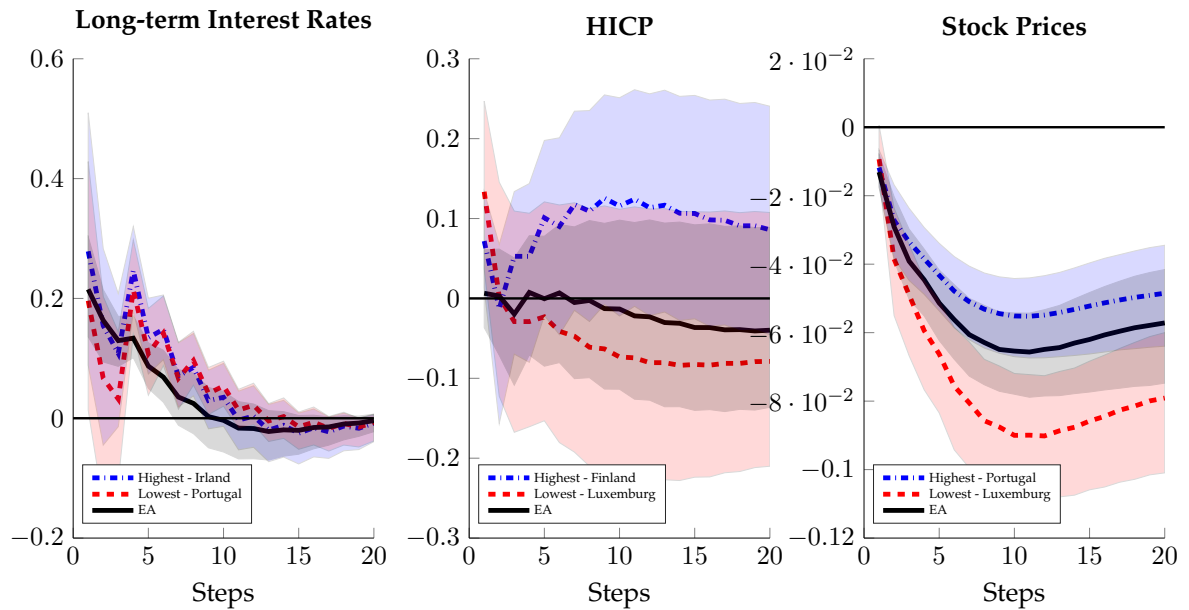

Figure 19: Highest/lowest percentage responses of selected prices to a 25bp contractionary policy shock across euro area member countries.

Chapter 2

Macroeconomic Stabilisation and Capital Controls

2.1 Introduction

For many decades, capital controls have caused controversy and led to intense debate amongst international policy makers. While the founders of the IMF, led by J.M. Keynes and H.D. White, fostered some scepticism of liberalised capital accounts and free international capital flows, the institutional view of the IMF soon turned in favour of free capital flows. The "Washington Consensus", a set of economic policy guidelines promoted by Washington-based institutions such as the IMF, the World Bank and the US Treasury Department in the second half of the 20th century, unambiguously spoke out in favour of free capital movements on a global scale. In doing so, they highlighted the potentially harmful aspects of capital controls: Similar to tariffs on goods trade, capital controls could change relative prices, giving one country a competitive advantage at the expense of another. Importantly, the assumption was that even in a second-best world, every step taken towards the liberalisation of financial markets would increase overall welfare. However, rapid capital flow reversals during the Latin American and South-East Asian crises in the 1980s and 90s fundamentally challenged this institutional view. Policy debate focussed on the ability of capital controls to prevent crisis phenomena, such as sudden stops and capital flight. The recent Great Recession and Global Financial Crisis have only served to reinforce debate about a more differentiated approach to capital controls, eventually even leading to the IMF's official adoption of a balanced view.[1]

Acknowledging the conflicting properties of capital controls—negative when coming at the cost of another country, positive when safeguarding financial stability—a host of recent research has set out to gain further insights about the mechanisms underlying capital controls. Heathcote

---

[1] see Ostry et al. (2010) and Ostry et al. (2012).

and Perri (2014) provide a particularly striking result that, even when leaving aside aspects of financial stability, capital controls can act as a macroeconomic stabiliser and increase risk sharing between countries. This finding, that in a second-best world closing some markets can be Pareto improving, has the potential to simplify the political process of implementing controls in the first place. Taking it as the starting point for my own research, I aim to deepen the understanding of how capital controls can be used to share risk internationally.

In summary, Heathcote and Perri (2014) show that capital controls can preserve risk sharing properties of terms of trade movements when pecuniary externalities would otherwise distort them. In their paper, the authors compare a model in which two countries can trade one non-contingent bond to a model of complete financial autarky. The latter is representative of a world with capital controls of the most extreme kind—a complete shut-down in the trade of all financial assets. The clarity of this approach allows the authors to showcase their main results. However, having such a blunt measure of capital controls also obscures the view to further insights. For this reason, I create a model that allows for a continuous spectrum of capital controls in the form of a tax on bond returns. The main questions in this context are: How large should taxes on international capital flows be to optimally insure against uncertainty? And how large are the additional welfare gains compared to the more blunt approach of shutting down asset trade completely?

The second aim of my research is to generalise the applicability of the framework used in Heathcote and Perri (2014) to a capital controls question. When examining the response of the terms of trade to a productivity shock, Heathcote and Perri (2014) restrict themselves to calibrations for the elasticity of substitution between domestically and foreign produced goods where the terms of trade *deteriorate* in response to a positive productivity shock. This can lead to what Bhagwati (1958) labelled "immiserising growth", when the deterioration in the terms of trade outweighs the gain in productivity. Consequently, capital controls that dampen the reaction of the terms of trade have the potential to increase welfare. Importantly, however, the reaction of the terms of trade to productivity shocks is discontinuous and non-monotonic in the trade elasticity. Corsetti, Dedola and Leduc (2008) show that a very low elasticity leads to the opposite reaction in the terms of trade. The logic goes that for very low elasticities, a deterioration in the terms of trade (i.e. a fall in the price of the relatively more abundant home good) cannot be an equilibrium solution, as demand from foreign consumers no longer matches supply. The only way to clear markets is to boost home income through a favourable move in the terms of trade. The question remains unanswered, however, whether the result that no trade in assets can be favourable to some trade in assets extends to this region.

In answering this question, I confront the paper's main computational challenge. As welfare differences between models are by nature small, it becomes necessary to find solutions using more

precise higher order approximations. In particular, I use third order approximations, which take various features into account that would be neglected in a linear approximation, most importantly a precautionary savings motive. Whilst the steady state of small open economies with incomplete asset markets depends on initial conditions, the effects of such higher order approximation can help to pin down steady-state values—in particular the bond position—and thus close the model. However, Andreasen et al. (2013) show that using third order approximations can also produce explosive sample paths which leads to the unconditional mean of the sample being different from the steady state. Their suggested solution is a systematic pruning of the state space and thus removing terms of orders higher than the approximation order. Doing so can in turn destroy the equilibrating effects of higher orders mentioned above. The exact workings of this interplay are subject of ongoing study, but current results suggest that an ad-hoc device as suggested in Schmitt-Grohe and Uribe (2003) can be used to equilibrate the model. In particular, combining a pruning approach with Uzawa preferences succeeds in overcoming both explosiveness issues of the sample path as well as equilibrating problems of the model itself.

The rest of the paper is structured as follows. Section 2.2 gives a detailed overview of the related literature with a particular focus on welfare analysis and capital controls. Section 2.3 presents the model. In Sections 2.4 and 2.5, I present the calibration as well as the computational approach. Results are presented in Section 2.6, whilst Section 2.7 concludes.

2.2   Related Literature

Whilst the literature on capital controls has expanded rapidly over recent years, there are varying rationales for including capital controls in a model. The two most common reasons for modelling capital controls are financial stability and terms of trade manipulations. This section will give a broad overview of both strands of the literature. Corresponding to the aim of this paper, however, the focus will be on the latter.

As mentioned above, the paper extends the work of Heathcote and Perri (2014). De Paoli and Lipinska (2013) present a very similar setup, but with a fundamentally different result. Whereas Heathcote and Perri find situations in which capital controls can increase insurance between countries, De Paoli and Lipinska state that restricting international capital movements "critically limits cross-border pooling of risk." This difference goes back to fundamental assumptions about capital in the economy. In Heathcote and Perri, endogenous capital accumulation magnifies pecuniary externalities. Critically, firms do not internalise the effect of their decisions on the terms of trade and overinvest after a positive shock to productivity, making the presence of capital accumulation in the model a crucial determinant of relative supply and demand. On the other hand, not featuring productive capital in their model, De Paoli and Lipinska (2013)

find that imposing capital controls will always only benefit one of the two countries in the model. Brunnermeier and Sannikov (2015) also look at terms of trade deteriorations when capital markets are open, but put the focus of their work on the effect of capital controls on financial stability. By introducing an adjustment cost to capital, their model exhibits an amplification mechanism where sudden stops and fire sales lead to sharp corrections in prices. Costinot et al. (2014) also focus on terms of trade movements, but look at them from the perspective of the "optimal tariff theory". The paper solves for the optimal capital flow tax a country has to implement in order to maximise the monopoly rent it can extract from other countries. The core of this argument can already be found in a discussion between Keynes (1929) and Ohlin (1929) in the context of German transfer payments after the first World War. The idea that a country can maximise its monopoly rent by introducing capital controls can also be found in a number of earlier contributions (see e.g. Obstfeld and Rogoff, 1996). Ostry et al. (2012) take a look at how one country's capital controls spill over into other countries. Addressing policy makers, they try to identify "rules of the road" to make sure multilateral considerations are taken into account whenever controls are employed. These spillover effects are estimated empirically in Lambert et al. (2011) and Forbes et al. (2012). A more general survey of empirical work on capital controls can be found in Magud et al. (2011).

Of the models focussing on financial stability, recent contributions include Bianchi and Mendoza (2013), where countries face default risk. Capital controls improve financial stability as individual agents do not internalise the effect of their actions on the government's incentives to default. In Bianchi (2011) and Korinek (2011) asset-price movements tighten an exogenously imposed collateral constraint, leading to inefficiencies (see also Caballero and Krishnamurty, 2004). In Martin and Taddei (2013), capital controls help to smooth informational frictions arising from private information. Another noteworthy strand of literature explores the benefits of capital controls in settings with price and wage rigidities. Amongst them, Farhi and Werning (2014a) find that, defying the Mundellian view, capital controls are desirable even when the exchange rate is flexible. In a different contribution, Farhi and Werning (2014b) present how ex-post contingent transfers can implement an efficient insurance arrangement in fiscal unions.

The main mechanism through which capital controls affect outcomes in this paper are movements in the terms of trade. Helpman and Razin (1978) were the first ones to bring forward the idea that these movements can provide an automatic insurance mechanism. The idea has been formalised in a seminal contribution by Cole and Obstfeld (1991), showing that under specific circumstances, in particular a unitary elasticity of substitution between traded goods, this insurance can be perfect.[2] Deviating from this case, Corsetti et al. (2008) show that terms

---

[2]The parameterisation under which insurance is perfect has subsequently become known as the "Cole-Obstfeld" case.

of trade movements can be very large when elasticities of substitution are low, giving rise to a potential role for policy intervention. Taking a closer look at cross-border insurance, the authors demonstrate that financial flows and relative prices can be either complements or substitutes in providing insurance. It is this interplay between financial flows and relative prices which will be central to understanding the mechanisms at play in this paper.

Finally, it is necessary to share a word on the computational literature that has been quintessential for this paper. Given the dimensionality of the state-space in New Open Economy Macroeconomic (NOEM) models, the parsimony owing to perturbation methods is particularly sought-after. Moreover, to attain accurate welfare calculations, and economically meaningful differences, I use a third-order approximation. Although higher-order approximations are straightforward to compute, they regularly generate explosive simulation paths arising from the additional fixed points that higher-order terms induce (Den Haan and De Wind, 2012). Kim, Kim, Schaumburg and Sims (2008) were amongst the first to propose a solution to this problem: A pruning approach—leaving out terms in the solution that have effects of higher order than the approximation—to high-order approximations for non-linear stochastic rational expectations models. In particular, they described an algorithm for calculating second-order approximations and show that explosive simulation paths are avoided. Based on their work, Andreasen et al. (2013) propose a generalised pruning method to calculate approximations above second order, exploring the econometric and macroeconomic implications of the method.[3] In this paper, I use the pruning algorithms of Andreasen et al. (2013). The welfare calculations presented below rely on simulations, making it paramount that simulations are stationary. To the best of my knowledge, this is the first application of pruning methods to the NOEM literature.

However, the application of pruning to NOEM models is complicated by another strand of literature emphasising computational issues with open economy models as well as questions of multiplicity. Schmitt-Grohe and Uribe (2003) show that a linearised small open economy model can have an undetermined bond position in equilibrium and propose tools to pin it down.[4] Schmitt-Grohe and Uribe (2003) advocate the use of 'stationarity-inducing devices', such as Uzawa-type preferences, to close such models. Heathcote and Perri (2014) circumvent this issue with a third-order approximation. In their model, a third-order approximation captures effects such as a precautionary savings motive that determines the models equilibrium. When applying a pruning approach, however, higher-order terms are removed and with them parts of equilibrating effects such as the precautionary savings motive that are needed to close the

---

[3]In related research, Lombardo and Uhlig (2014) provide a theory of pruning and formulas for pruning of any order and relate the method to results described by Judd (1998) on perturbing dynamical systems.

[4]Bodenstein (2011) demonstrates that large open economies can suffer from multiple steady states (even in financial autarky) and in turn proposes methods to eliminate all but one equilibrium. Bodenstein (2011) emphasises that this issue is distinct from the non-stationarity discussed in Schmitt-Grohe and Uribe (2003). Here I circumvent Bodenstein's critique by calibration and steady state definition.

model. To solve both problems here—explosiveness of simulations and non-stationarity—I use Uzawa-type preferences alongside pruning.

2.3 The Model

The general setup follows the model in Heathcote and Perri (2014), which in turn is an adaptation of the canonical Backus, Kehoe and Kydland (1992) two-country, two-good model. Each country $i = 1, 2$ is characterised by two types of profit-maximising firms, producing traded intermediary goods and non-traded final goods, the latter used for consumption and investment. Identical, infinitely-lived households own the companies, consume final goods, provide labour, rent out capital and make investment decisions. Every period t, one state of the world $s_t \in S$ is realised. The history of realised states up to and including period t is expressed by $s^t \equiv (s_1, ..., s_t)$.

2.3.1 Preferences and Technologies

The representative agent in each country i maximises expected lifetime utility given by

$$E_0 \sum_{t=0}^{\infty} [\beta_i(s^t)]^t U(c_i(s^t), n_i(s^t)), \tag{2.1}$$

where $\beta_i(s^t)$ denotes the discount factor for future periods, defined, unless otherwise stated, to be $\beta \in (0, 1)$. $\pi(s^t)$ is the probability at date 0 of history s^t, and c_{it} and n_{it} are period consumption and labour, respectively, in country i. I consider a standard utility function which is separable in consumption and hours worked:

$$U(c_i, n_i) = \frac{c_i^{1-\gamma}}{1-\gamma} - \phi \frac{n_i^{1+\frac{1}{\epsilon}}}{1+\frac{1}{\epsilon}}, \tag{2.2}$$

where γ determines the degree of risk aversion, while ϵ is the Frisch elasticity of labour supply. Under the baseline calibration of $\gamma, \epsilon = 1$, the utility function becomes separable between the log of consumption and a term involving hours worked:

$$U(c_i, n_i) = \log c_{it} - \phi \frac{n_{it}^2}{2}. \tag{2.3}$$

Intermediate goods firms produce country specific goods a (in Country 1) and b (in Country 2), using a Cobb-Douglas production technology of the form

$$F_i(z_i, k_i, n_i) = \exp(z_i) k_i^{\theta} (L_i n_i)^{1-\theta}, \tag{2.4}$$

where z_i is an exogenous productivity shock, k_i is accumulated country-specific capital employed,

L_i is the labour endowment of the country (which I interchangeably also interpret as the size of the country), and θ defines the share of capital in the production. Goods a and b are the only traded goods in the model.

The productivity shock evolves without spillovers according to the symmetric autoregressive process

$$z_i(s^t) = \rho z_i(s^{t-1}) + \epsilon_i(s^t) \tag{2.5}$$

$$\begin{bmatrix} \epsilon_1(s^t) \\ \epsilon_2(s^t) \end{bmatrix} \sim N \left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \sigma_\epsilon^2 \begin{pmatrix} 1 & Corr_{\epsilon_1, \epsilon_2} \\ Corr_{\epsilon_1, \epsilon_2} & 1 \end{pmatrix} \right]. \tag{2.6}$$

Final goods in each country are produced using a combination of intermediate goods a and b as inputs, according to the CES production functions

$$G_1(a_1, b_1) = \left[\omega a_1^{\frac{\sigma-1}{\sigma}} + (1-\omega) b_1^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \tag{2.7}$$

$$G_2(a_2, b_2) = \left[(1-\omega) a_2^{\frac{\sigma-1}{\sigma}} + \omega b_2^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \tag{2.8}$$

for Country 1 and 2, respectively, with a_i and b_i denoting the quantity of goods a and b employed in country i. The coefficient $\omega \in \{0, 1\}$ determines the ratio between home and foreign intermediary goods used in the production of each country's final good. A value of $\omega > 0.5$ results in home bias, whereas a value below 0.5 produces an anti-home or foreign bias. The coefficient σ determines the elasticity of substitution between home and foreign-produced intermediate goods. In the limit as $\sigma \to 1$, the production function takes a Cobb-Douglas form. The final good is used for both consumption and investment, augmenting the capital stock in the following way:

$$k_i(s^t) = (1-\delta)k_i(s^{t-1}) + x_i(s^t), \tag{2.9}$$

where $x_i(s^t)$ is the amount of the country i final good used for investment in country i, and δ is the rate of depreciation.

The resulting resource constraints in the model are

$$F(z_1(s^t), k_1(s^{t-1}), n_1(s^t)) = a_1(s^t) + a_2(s^t) \tag{2.10}$$

$$F(z_2(s^t), k_2(s^{t-1}), n_2(s^t)) = b_1(s^t) + b_2(s^t) \tag{2.11}$$

for intermediate goods, and

$$G_i(a_i(s^t), b_i(s^t)) = c_i(s^t) + x_i(s^t) \tag{2.12}$$

for final goods, where $i = 1, 2$.

2.3.2 Firm Problems

Intermediate goods-producing firms hire labour and capital from households at market rates $w_i(s^t)$ and $r_i(s^t)$, respectively, in units of the final good.[5] Their maximisation problems in countries 1 and 2 are, respectively:

$$\max_{n_1(s^t), k_1(s^{t-1})} \{q_1^a(s^t)F(z_1(s^t), k_1(s^{t-1}), n_1(s^t)) - w_1(s^t)n_1(s^t) - r_1(s^t)k_1(s^{t-1})\} \qquad (2.13)$$

$$\max_{n_2(s^t), k_2(s^{t-1})} \{q_2^b(s^t)F(z_2(s^t), k_2(s^{t-1}), n_2(s^t)) - w_2(s^t)n_2(s^t) - r_2(s^t)k_2(s^{t-1})\}, \qquad (2.14)$$

where $q_i^a(s^t)$ and $q_i^b(s^t)$ are the prices at which intermediate goods are traded in country i.

Final goods-producing firms then buy inputs $a_i(s^t)$ and $b_i(s^t)$ on the home and foreign intermediate goods markets to maximise their static optimisation problem:

$$\max_{a_i(s^t), b_i(s^t)} \{G_i(a_i(s^t), b_i(s^t)) - q_i^a(s^t)a_i(s^t) - q_i^b(s^t)b_i(s^t)\}. \qquad (2.15)$$

2.3.3 Terms of Trade and the Real Exchange Rate

As described above, one of the central stabilising properties of the model will rely on the movement in the terms of trade $p(s^t)$. They are defined as the price of the foreign good relative to the home good. In equilibrium, this corresponds to the marginal rate of transformation between goods a and b in country i:

$$p(s^t) = \frac{q_1^b(s^t)}{q_1^a(s^t)} = \frac{q_2^b(s^t)}{q_2^a(s^t)} = \frac{\partial G_i(a_i(s^t), b_i(s^t))/\partial b_i(s^t)}{\partial G_i(a_i(s^t), b_i(s^t))/\partial a_i(s^t)}. \qquad (2.16)$$

By this definition, an *increase* in the terms of trade is equal to a *deterioration* from the point of view of Country 1.

The real exchange rate $rx(s^t)$ is defined as the ratio of prices paid for the same intermediate good in Country 1 and Country 2:

$$rx(s^t) = \frac{q_1^a(s^t)}{q_2^a(s^t)} = \frac{q_1^b(s^t)}{q_2^b(s^t)}. \qquad (2.17)$$

By this definition, too, an *increase* in the real exchange rate is equal to a *deterioration* from the point of view of Country 1.

---

[5]In a number of closely related setups, intermediate goods-producing firms own the capital and make investment decisions themselves. However, for the intertemporal decision between paying out dividends and investing in the capital stock, I assume that firms use the discount factor of the representative household. Modelling ownership of capital directly with the households hence does not change the analytical result, whilst improving the clarity of the exposition.

2.3.4   Household Problem and Capital Controls

Households maximise lifetime utility subject to a budget constraint, which leaves them with decisions about consumption, supplying labour to intermediary goods-producing firms, investing in productive capital and, if possible, investing in financial assets. As the possibility to invest in financial assets constitutes the main interest of this research, I will focus on comparing different setups that allow agents to make investment decisions in a more or less constrained way.

Complete Markets   The benchmark case is given by a complete markets (CM) scenario, where agents can invest in a full set of state-contingent securities at price $P(s^t, s_{t+1})$, equal to the number of possible states of the world. To achieve symmetry in the results, each security pays out half a unit of the domestic and half a unit of the foreign consumption good.[6] The corresponding budget constraints for countries 1 and 2 are given by

$$c_1(s^t) + \sum_{s_{t+1}} P(s^t, s_{t+1}) B_1(s^t, s_{t+1}) + x_1(s^t) =$$
$$w_1(s^t) n_1(s^t) L_1 + r_1(s^t) k_1(s^{t-1}) + B_1(s^{t-1}, s_t)\Big(\frac{1}{2} + \frac{1}{2} rx(s^t)\Big) \quad (2.18)$$

$$c_2(s^t) + \sum_{s_{t+1}} P(s^t, s_{t+1}) \frac{B_2(s^t, s_{t+1})}{rx(s^t)} + x_2(s^t) =$$
$$w_2(s^t) n_2(s^t) L_2 + r_2(s^t) k_2(s^{t-1}) + B_2(s^{t-1}, s_t)\Big(\frac{1}{2}\frac{1}{rx(s^t)} + \frac{1}{2}\Big). \quad (2.19)$$

Bond Economy   Secondly, I will look at a bond economy (BE) where agents have the option to invest in only one international non-contingent bond. To preserve symmetry, the bond pays out half a unit of the domestic and half a unit of the foreign consumption good. A particular interest of the paper lies in the introduction of a moderate form of capital controls, where this bond can still be traded, but facing a tax τ_i. To this end, I introduce a BE model with taxes that nests a pure BE at $\tau = 0$. The budget constraints for this case are

$$c_1(s^t) + P(s^t) B_1(s^t) + x_1(s^t) =$$
$$w_1(s^t) n_1(s^t) L_1 + r_1(s^t) k_1(s^{t-1}) + (1 + \tau_1(s^t)) B_1(s^{t-1})\Big(\frac{1}{2} + \frac{1}{2} rx(s^t)\Big) - Tr_1 \quad (2.20)$$

---

[6]Although seemingly unrealistic, the assumption about bond payoffs does not change the underlying mechanisms. The bond could be defined as paying out one unit of the domestic good or, alternatively, as paying out one unit of the foreign good without altering the extent to which countries can share risks. The only difference lies in the property that symmetry is perfectly preserved under the specified bond. This in turn makes it possible to compare welfare more easily across countries and models.

$$c_2(s^t) + P(s^t)\frac{B_2(s^t, s_{t+1})}{rx(s^t)} + x_2(s^t) =$$

$$w_2(s^t)n_2(s^t)L_2 + r_2(s^t)k_2(s^{t-1}) + (1 + \tau_2(s^t))B_2(s^{t-1})\left(\frac{1}{2}\frac{1}{rx(s^t)} + \frac{1}{2}\right) - Tr_2, \quad (2.21)$$

where Tr_i is a lump-sum transfer of the size

$$Tr_i = \begin{cases} \tau_1(s^t)B_1(s^{t-1})\left(\frac{1}{2} + \frac{1}{2}rx(s^t)\right) & , i = 1 \\ \tau_2(s^t)B_2(s^{t-1})\left(\frac{1}{2}\frac{1}{rx(s^t)} + \frac{1}{2}\right) & , i = 2. \end{cases} \quad (2.22)$$

Financial Autarky As the limiting case for the scenario of extremely high capital controls which lead to a complete shut-down of asset trade, I will finally take a look at a financial autarky (FA) model. The households' budget constraints for this case reduce to

$$c_i(s^t) + x_i(s^t) = w_i(s^t)n_i(s^t)L_i + r_i(s^t)k_i(s^{t-1}) \qquad i = 1, 2. \quad (2.23)$$

One interesting aspect of the numerical solution will lie in replicating the FA model through the imposition of prohibitively high capital controls.

There are four different constellations through which a capital controls tax can affect a country's agents:

- $\tau_i > 0$, $B_i > 0$: Subsidy on international lending. The country is a net-lender and accumulates bonds. With a positive tax, the return on the bond position gets larger, making agents wanting to increase their bond position further.

- $\tau_i > 0$, $B_i < 0$: Tax on international borrowing. As the country is a net-borrower, having a positive tax on bond returns makes it more expensive to borrow. Agents will want to reduce their bond holdings.

- $\tau_i < 0$, $B_i > 0$: Tax on international lending. The country is a net-lender, but the tax on bonds reduces returns. Compared to no tax, agents will want to hold less bonds.

- $\tau_i < 0$, $B_i < 0$: Subsidy on international borrowing. The negative tax reduces the costs of borrowing money. Agents will want to borrow more.

It becomes clear that a capital controls tax seeking to reduce trade in bonds has to be of the opposite sign as a country's bond position, whereas a policy to increase borrowing and lending needs both tax and bond position to be of the same sign. In other words, even taking the simplifying assumption of constant taxes in absolute value,[7] the sign of the tax has to change

---

[7]Simplifying taxes by making them constant over time is in fact a significant step in taking the model closer to reality. As documented in Eichengreen and Rose (2014), capital controls are usually highly static, corresponding to a constant, rather than a constantly adjusted tax.

depending on the bond position of the country. This presents the computational challenge of solving the model despite the obvious break in the continuity of the functions of τ_i. To solve this problem, I define

$$\tau_i(s^t) = \bar{\tau} \cdot (-1)B_i(s^t). \tag{2.24}$$

In this way, the function is again differentiable and a perturbation approach produces a meaningful result. This setup has the additional welcome property that excessive asset trade is punished more the higher the already built-up bond position.

2.3.5 Definition of Equilibrium

An equilibrium is a set of prices $w_i(s^t)$, $r_i(s^t)$, $q_i^a(s^t)$, $q_i^b(s^t)$, $p(s^t)$ and $rx(s^t)$ for all s^t and all $t > 0$, such that, given an exogenous capital control $\bar{\tau}$

1. all agents solve their optimal consumption, labour, investment and portfolio choice problems, subject to budget constraints and

2. the goods market clearing conditions (10)-(12) hold, as do the asset market clearing conditions, i.e.

$$B_1(s^t) + B_2(s^t) = 0 \tag{2.25}$$

$$B_1(s^t, s_{t+1}) + B_2(s^t, s_{t+1}) = 0 \qquad \forall s_{t+1} \in S \tag{2.26}$$

for the bond economies and complete markets models, respectively.

2.4 Calibration

For the baseline scenario, I choose a calibration roughly reflecting properties of an emerging market economy. The discount factor β is set to 0.99, as is standard for a quarterly model. The weight on labour ϕ is calibrated to 7.3, so the hours worked are equal to 1/3 in steady state. As mentioned above, γ, the curvature of the utility function, is set to 1, giving utility a logarithmic form. Moreover, the Frisch elasticity ϵ is set to 1 as well. Following Heathcote and Perri (2013), I set θ, the share of GDP going to capital, to 36 percent and the depreciation rate δ to 0.015.

In the baseline scenario, the elasticity of substitution between home and foreign intermediate inputs, σ, is set to 1. Consequently, the technology of final goods producing firms becomes Cobb-Douglas. The import share, which is negatively related to the home-bias parameter ω, is set to $is = 0.25$, indicating a home bias in the production of final goods.[8] This corresponds to values found in empirical studies (Lewis, 1999).

---

[8]In steady state, the home-bias parameter and import share are related by $\left(\frac{\omega}{1-\omega}\right)^{\sigma} = \frac{1-is}{is}$.

Table 2.1: Parameter Values

Preferences			
	Discount factor	$\beta = 0.99$	
	Uzawa factor	$\alpha = 0$ (alternatively 0.01)	
	Weight on labour	$\phi = 7.3$	
	Curvature	$\gamma = 1$	
	Frisch Elasticity	$\epsilon = 1$	
Country size			
	Country 1	$L_1 = 1$ (alternatively 1.5 and 2)	
	Country 2	$L_2 = 1$	
Technology			
	Capital's share	$\theta = 0.36$	
	Depreciation rate	$\delta = 0.015$	
	Elasticity of substitution	$\sigma \in \{0.1, 0.2, 0.3, 0.4, 0.5, 1, 1.5, 2\}$	
	Import share	$is \in \{0.15, 0.25\}$	
Productivity process			
	Persistence	$\rho = 0.995$	
	Variance	$\sigma_\epsilon^2 = 0.02$ (alternatively 0.07)	
	Correlation	$Corr_{\epsilon_1, \epsilon_2} = 0.38$	

The calibration for the productivity process is identical to Heathcote and Perri (2014). The persistence of shocks is set to $\rho = 0.995$, the variance of innovations is set to $\sigma_\epsilon^2 = 0.02$ and the correlation of innovations $Corr_{\epsilon_1, \epsilon_2} = 0.38$. As Heathcote and Perri point out, the values for the variation and correlation of shocks are consistent with estimates for developing economies, as shown for example in Neumeyer and Perri (2005).

Moving away from the baseline calibration, I will analyse the effect of varying both the import share and the elasticity of substitution between inputs in the final goods production process. On the import share, the extended calibrations will look at $is = 0.25$ and $is = 0.15$. Whereas an import share of 25% is representative of emerging economies such as China, Russia, Indonesia and Bangladesh, an import share of around 15% is found in countries such as Nigeria, Argentina and Brazil (World Bank, 2015). Both values hence help to simulate the model close to real-world emerging markets.

The bulk of the analysis will focus on two identically large countries ($L_i = 1$ for $i = 1, 2$) to showcase underlying mechanisms as clearly as possible. For robustness checks on the extended version of the model, I will first calibrate the size of Country 1 to $L_i = 1.5$, whilst normalising Country 2's size to 1. Finally, I will specify the size of Country 1 to be twice the size of Country 2.

A parameter that caused controversy in the literature is the elasticity of substitution for traded goods. Whilst this paper does not want to take a stance as to which particular value might be "correct", it seems crucial to a well-founded analysis to take into account all values that are likely to have external validity. For this purpose, I investigate a parameter span between

$\sigma = 0.1$ and 2, as supported, amongst others, by Hooper et al. (2000), who reports a short-run trade elasticity of 0.6 for the US and values between 0 and 0.6 for the remaining G7 countries, or Taylor (1993), who finds a short-run trade elasticity of 0.22. As shown in Corsetti et al. (2008), in response to shocks, the terms of trade will exhibit small positive and negative reactions at the upper and lower end of the calibration, respectively, whilst featuring very large reactions, both positive and negative, for intermediate values. Given that the model dynamics so fundamentally change with the parameterisation of σ, it seems crucial both from a theoretical as well as policy perspective to extend the analysis of capital controls in that respect. To my knowledge, no earlier paper has done so.

2.5 Computation

In the complete markets (CM) benchmark case, the welfare theorems and assumed functional forms allow to solve for the planning problem instead of the decentralised equilibrium. In particular, this is expressed in an intertemporal risk sharing condition for a and b (see Appendix A for all equilibrium conditions). For the financial autarky (FA) and bond economy (BE) setting, it is necessary to solve for the decentralised equilibrium and prices explicitly. In the former case the modified budget constraints together with Walras' Law will ensure that markets clear. For the BE, however, an additional intertemporal Euler equation for bonds needs to be introduced.

As differences in welfare between the models are by nature very small, it becomes fundamental to increase the precision of the result through a higher-order approximation. More specifically, I solve the baseline calibration by performing 3rd order local approximations around the non-stochastic steady state, using the software platform DYNARE. The third order approximation captures not only the effect of uncertainty on optimal choice, but also how this effect changes with the level of state variables. Given the assumptions made about preferences, agents in the model exhibit a precautionary savings motive. However, as the agents' wealth increases with their bond position, the motive for further precautionary savings is reduced.[9] Under baseline parameter values, these effects are strong enough to close the model and consequently produce a stable bond position. Moving away from the baseline calibration, however, some terms in the approximation become explosive and overpower the stabilising effects mentioned above. I implement two solutions to overcome this problem. Firstly, as suggested in Schmitt-Grohe and Uribe (2003) and Bodenstein (2011), the model is augmented with endogenous discount factors, also known as *Uzawa preferences*, after Uzawa (1968). Specifically, the discount factors will be defined as

$$\beta_i(s^t) = \bar{\beta} \cdot c_i(s^t)^{-\alpha}, \tag{2.27}$$

---

[9]This corresponds to a decreasing coefficient of absolute risk aversion, defined as $R_A = -\frac{u''(c)}{u'(c)}$. This is given in the case of logarithmic utility, as $R_A = \frac{1}{c}$.

where $\alpha > 0$ determines the weight of consumption in the discount factor and is calibrated to 0.01, whilst $\bar{\beta}$ is calibrated so that β_i equals 0.99 in steady state. In effect, the more an agent consumes (i.e. the richer she is), the more impatient she becomes, driving down her savings (i.e. her bond position). As pointed out in Schmitt-Grohe and Uribe (2003), this modification makes the steady state independent of initial conditions while at the same time preserving comovements of macroeconomic aggregates from the baseline model.

Secondly, I adopt the solution proposed by Andreasen et al. (2013) of systematically pruning the state space. Put simply, this means removing terms in the solution that have higher-order effects than the approximation order. Doing so ensures that simulation paths do not explode. That said, pruning also removes higher order terms that help to equilibrate the model. It is in connection to this issue that Uzawa preferences restore equilibrium by providing a substitute for equilibrating terms that were removed under pruning.

This first application of pruning in an open economy macroeconomics setting has shown to be very promising. At the same time, a host of questions about the exact workings at play were thrown up, which soon outgrew the scope of this paper and have become the subject of separate ongoing study.

To compare welfare across models with and without capital controls, I employ a consumption equivalent measure similar to Schmitt-Grohe and Uribe (2007). In other words, I ask what percentage of permanent increase in consumption under one regime would bring welfare to the same level as in the other. I denote the welfare associated with a bond economy without capital controls by the expected lifetime utility in period 0, when the economy is in it's non-stochastic steady state:

$$V_{i0}^{BE} = E_0 \sum_{t=0}^{\infty} \beta^t U(c_{it}^{BE}, n_{it}^{BE}). \tag{2.28}$$

Similarly, defining the welfare of a model with capital controls by

$$V_{i0}^{CC} = E_0 \sum_{t=0}^{\infty} \beta^t U(c_{it}^{CC}, n_{it}^{CC}), \tag{2.29}$$

we can compare the different results of welfare using a Lucas-style measure for consumption equivalence. I denote with ω_i the welfare gain of moving from the unrestricted bond economy to a system with capital controls. Defining ω_i as a fraction of consumption, we can write lifetime utility under capital controls (CC) as

$$V_{i0}^{CC} = E_0 \sum_{t=0}^{\infty} \beta^t U((1 + \omega_i)c_{it}^{BE}, n_{it}^{BE}). \tag{2.30}$$

Using the functional form of utility, we get

$$
\begin{aligned}
V_{i0}^{CC} &= E_0 \sum_{t=0}^{\infty} \beta^t \left[ln((1+\omega_i)c_{it}^{BE}) - \phi \frac{(n_{it}^{BE})^2}{2} \right] \\
&= \frac{ln(1+\omega_i)}{1-\beta} + E_0 \sum_{t=0}^{\infty} \beta^t \left[ln(c_{it}^{BE}) - \phi \frac{(n_{it}^{BE})^2}{2} \right] \\
&= \frac{ln(1+\omega_i)}{1-\beta} + V_{i0}^{BE}.
\end{aligned}
\tag{2.31}
$$

Solving for ω_i yields

$$
\omega_i = exp[(1-\beta)(V_{i0}^{CC} - V_{i0}^{BE})] - 1.
\tag{2.32}
$$

It is important to emphasise that for this approach to deliver a sensible result, both welfare measures have to be computed starting from the non-stochastic steady state, with identical stocks of capital and an identical bond position.

2.6 Results

The imposition of capital controls has a significant effect on allocative efficiency and risk sharing between countries. As capital wants to be employed where its marginal product is highest, restricting capital movement can only have a negative effect on allocative efficiency. Risk sharing, on the other hand, can go both ways. The main interest of this section will be to see whether a positive effect from risk sharing can ever outweigh the loss from a less efficient allocation.

2.6.1 Symmetric Countries

Welfare gains from imposing Financial Autarky				
	$\sigma = 0.5$	$\sigma = 1$	$\sigma = 1.5$	$\sigma = 2$
$is = 0.25$	0.060	-0.006	-0.018	-0.029
$is = 0.15$	0.108	-0.016	-0.014	-0.033

Table 2.2: Welfare gains calculated relative to bond economy, as % of consumption; un-pruned solution, standard preferences. is = import share, σ = elasticity of substitution between traded goods.

Table 2.2 shows the welfare gains of moving from the bond economy to complete financial autarky, results that are in line with those in Heathcote and Perri (2014). As can be seen, for the benchmark calibration with a unitary elasticity of substitution ($\sigma = 1$) and a 25% import share, imposing financial autarky leads to a loss in welfare equal to 0.006% of permanent consumption. What is the underlying mechanism? After a positive shock to productivity in Country 1, firms in the country want to expand their production capacity to make use of the higher productivity. Under the bond economy, they borrow from abroad to drive up investment

(see figure 2.1). As output for good a goes up, its price goes down and the terms of trade deteriorate (by definition, an increase in the terms of trade is bad for Country 1, but in favour of Country 2). Overall, demand shifts from good b to good a, allowing Country 2 to produce relatively less of good b. At the same time, the increase in the price of b increases the country's income, shifting some of the benefits from the increase in Country 1's productivity to Country 2. Overall, the possibility to borrow from abroad makes it possible to increase production where productivity is highest - an improvement in allocative efficiency - whilst the movement in the terms of trade distributes the gains across countries.
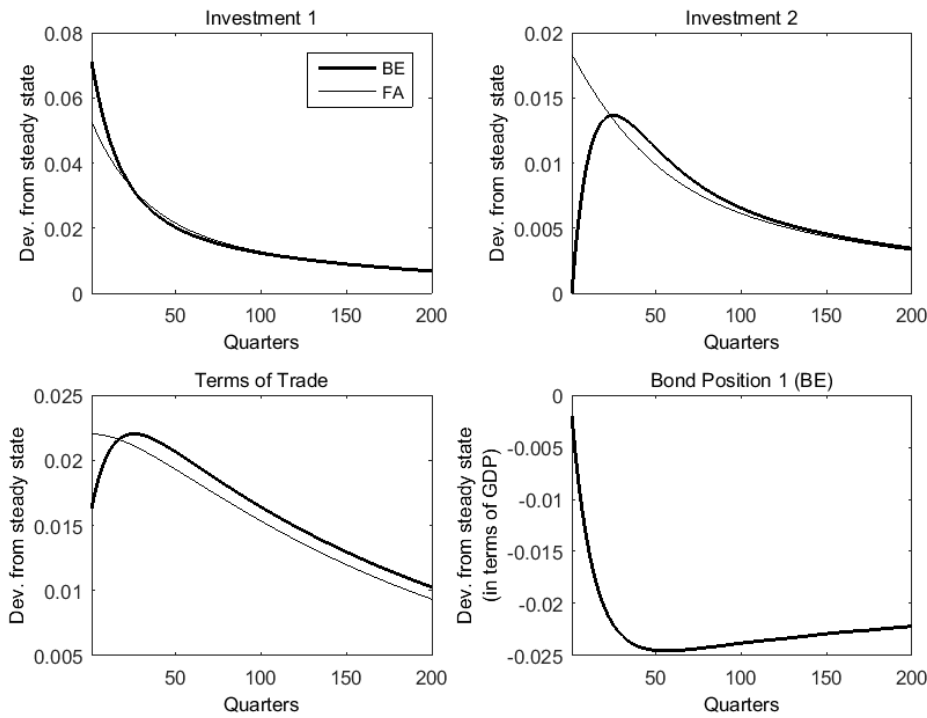

Figure 2.1: Impulse responses to a 1 standard deviation shock in Country 1's productivity. $\sigma = 1$, import share = 25%.

By imposing financial autarky, these positive effects are dampened, leading to a Pareto inferior outcome. Investment still increases in Country 1, but not as much as would be efficient. At the same time, the movement in the terms of trade is dampened, leading to a loss of cross-border insurance.

It is worth noting that this loss of cross-border insurance would benefit the country that has experienced the positive shock, as it could retain a larger share of the rise in output. This explains why a country might want to unilaterally impose capital controls after a positive shock has occurred. In other words, through the imposition of capital controls, domestic firms wouldn't be able to borrow from abroad and the build-up in production capacity would be less extensive. Overall, the country could sell its products at a higher price and, like a monopolist, maximise its rent. From an ex-ante perspective, however, no country knows whether it will be hit by

a positive or a negative shock and, at least for $\sigma = 1$, would prefer to keep capital markets unrestrained.

Pareto Improving Capital Controls

The result of the welfare comparison is fundamentally different when looking at the calibration with a lower elasticity of substitution of $\sigma = 0.5$. In fact, both countries would prefer to agree on capital controls of the strictest form before any shocks occur, with welfare improving by 0.06% of consumption over an agent's lifetime for the case of a 25% import share, and 0.108% at a 15% import share. From a welfare perspective, the loss in allocative efficiency is outweighed by an even greater gain in risk sharing. It is worth taking a closer look at the exact mechanism behind this.
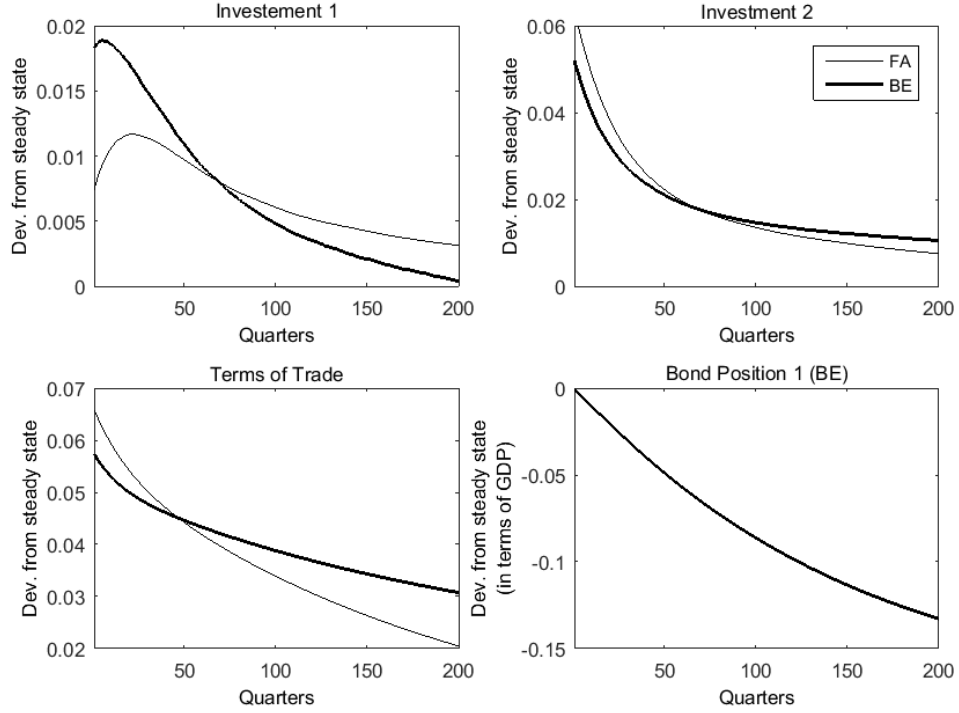

Figure 2.2: Impulse responses to a 1 standard deviation shock in Country 1's productivity. $\sigma = 0.5$, import share $= 25\%$.

As in the case above, a positive shock in productivity leads to an increase in investment (see figure 2.2). In a bond economy, agents borrow from abroad to invest even more. They do so without internalising that collectively this has an effect on prices. In fact, due to the low elasticity, the terms of trade worsen so much for the country with the positive shock that it ends up with lower welfare than before the shock. Without home bias, the agents from the other country would switch to the now cheaper good and prop prices back up. With home bias, as specified, however, agents in the other country do not buy much more of the cheaper good despite the fall in prices. If we associate the movements in the terms of trade with insurance,

this calibration would exemplify a typical case of overinsurance. The movements in the terms of trade are higher under the bond economy than under financial autarky. By imposing capital controls, the ability of agents to borrow from abroad is limited and overinvestment dampened. The top left panel of figure 2.2 shows investment in Country 1 cut roughly by half in the first periods after the shock. The resulting movement in the terms of trade is curbed[10] which in turn reduces the volatility of consumption and utility. How decisive movements in the terms of trade are for the position of a country can also be seen by looking at the absolute level of investment in the two countries. As depicted in the two top panels in figure 2.2, the movement in the terms of trade put Country 2 in such a strong position that it can afford a higher level of investment than Country 1, despite the latter having undergone a positive shock.

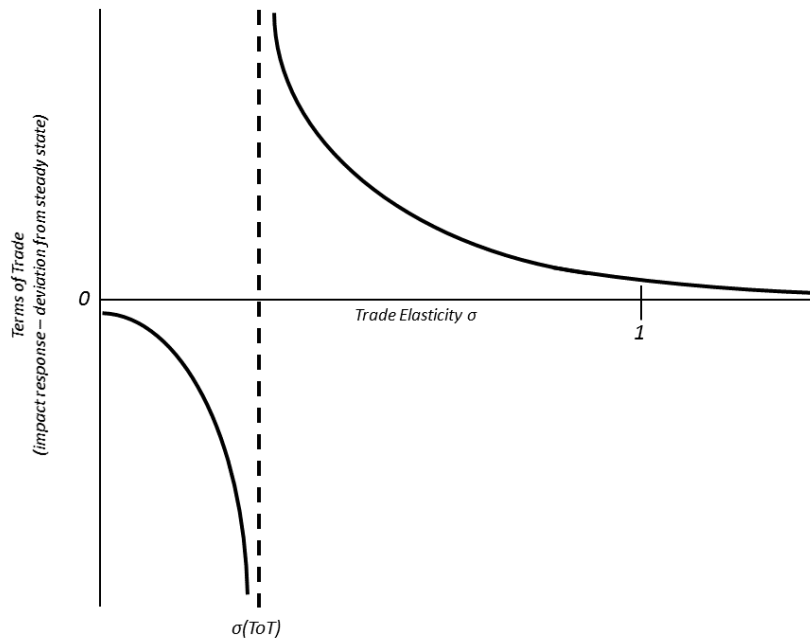**Discontinuity of Terms of Trade Reaction**

Figure 2.3: Schematic terms of trade response to a productivity shock in Country 1, as a function of trade elasticity σ in financial autarky.

In the cases mentioned above, a positive shock in productivity is followed by a deterioration in the terms of trade. Importantly however, this is not a general result, as demonstrated in Corsetti et al. (2008). Figure 2.3 schematically shows how the impact response of the terms of trade to a shock changes for different values of the trade elasticity. Specifically, it shows how the terms of trade reaction becomes continually larger as the trade elasticity becomes smaller, until, at point $\sigma(ToT)$, the response is non-monotonic, switches sign and approaches 0 from below (see Appendix B for an approach for finding $\sigma(ToT)$ analytically). The logic is that for

---

[10]Strictly speaking, the integral under the impulse response function for the terms of trade is smaller under FA than BE (see bottom left panel in figure 2.2).

very low elasticities, a rise in the terms of trade (i.e. a fall in the price of the relatively more abundant good) can no longer be an equilibrium solution, as demand from foreign consumers does not rise enough to match supply. For markets to clear, domestic demand must absorb the excess supply. This only happens if domestic income rises enough to do so. The only way for this to happen is through a favourable move in the terms of trade. Consequently, the terms of trade response switches sign. How does this affect the results?

	Welfare gains from imposing Financial Autarky				
	$\sigma = 0.1$	$\sigma = 0.2$	$\sigma = 0.3$	$\sigma = 0.4$	$\sigma = 0.5$
$is = 0.25$	3.714	no sol.	3.693	0.000	-0.038
$is = 0.15$	0.120	0.410	4.272	no sol.	0.0413

Table 2.3: Welfare gains calculated relative to bond economy, as % of consumption; pruned solution, Uzawa preferences. is = import share, σ = elasticity of substitution between traded goods.

As mentioned above, the explosive behaviour of the terms of trade make it impossible to find a stable solution for low elasticities in the textbook model. However, combining a pruning algorithm and Uzawa preferences allows to contain the explosive behaviour to some degree, giving solutions for a much wider range of parameters. Looking at the welfare comparisons in Table 2.3, a number of things are of note.

Firstly, the changed model and computational technique gives slightly different results from a quantitative point of view. At $\sigma = 0.5$ and $is = 0.25$, for example, this model shows a welfare loss of imposing capital controls, whereas the textbook model produced small gains. That said the underlying mechanisms and model behaviour are unchanged.

Secondly, the larger the absolute movement in the terms of trade, the larger the potential welfare gains from the imposition of capital controls. Moving from right to left (both in Figure 2.3 and Table 2.3) welfare gains increase as the terms of trade reaction gets larger. At $\sigma(ToT)$, the model does not have a solution. For an import share of 25%, this is the case in the neighbourhood of $\sigma = 0.2$, whereas it is around $\sigma = 0.4$ for an import share of 15%. Further to the left, where the reaction of the terms of trade is negative, welfare gains are large, but diminish as the terms of trade reaction approaches zero. Interestingly, however, it remains positive until the lowest level for which the model can be solved computationally.

To better understand the mechanisms at play, I will take a closer look at a calibration below $\sigma(ToT)$, namely $\sigma = 0.1, is = 0.25$. Figure 2.5 shows that under a bond economy, a positive productivity shock in Country 1 makes the country richer through an improvement in the terms of trade. The flip side is that Country 2 undergoes significant cuts in its income. In fact, the drop for Country 2 is so large that in the bond economy agents start borrowing funds from Country 1 by selling bonds to dampen the fall in income. By doing so, they pile up a large negative bond position, worsening the country's long-term prospects. More importantly, Country 1 can only
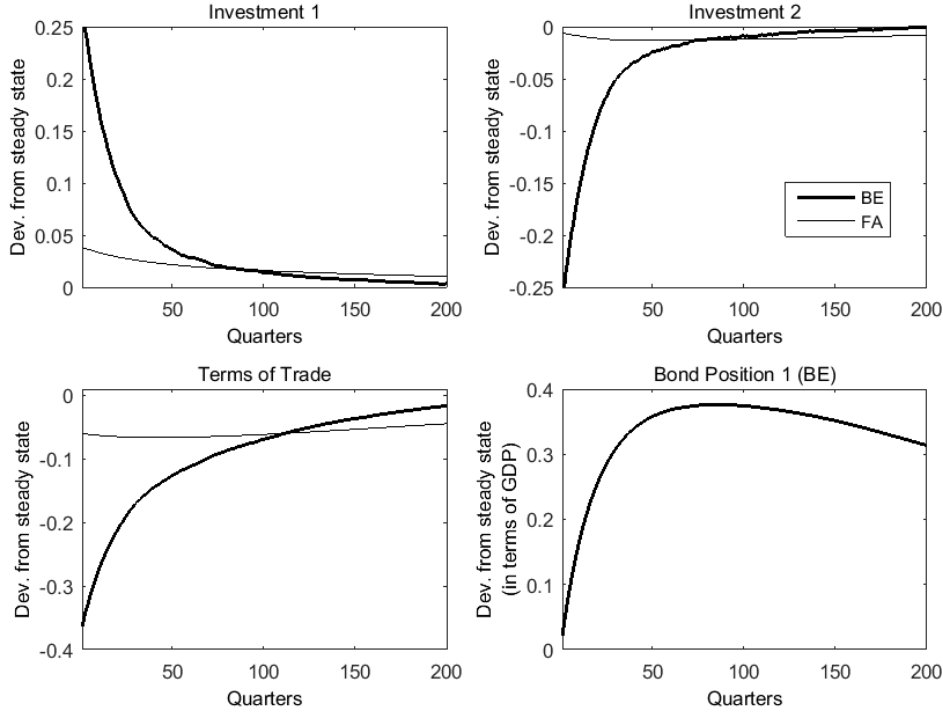
Figure 2.4: Impulse responses to a 1 standard deviation shock in Country 1's productivity. $\sigma = 0.1$, import share $= 25\%$.

become a lender if its income rises even more, i.e. if the terms of trade appreciate even more. Hence, far from helping to share risks across borders, asset trade in fact amplifies them. Moving to financial autarky, the terms of trade remain contained, and so does investment, output and relative consumption (see figure 2.5). Simulating the economy for a length of one million periods shows that the consumption variance under FA is only half the variance under BE. The variance of the return on capital investments is four times larger under BE than FA whilst the variance in the hours worked is ten times as large under BE than FA.

Looking at this from a slightly different angle, the question arises why Country 1 doesn't borrow from abroad to make more use of its increased productivity? Critically, if Country 1 did indeed borrow, i.e. incur a current account deficit, the terms of trade would not have to appreciate as much as they do. The boost in income from a smaller appreciation, together with funds from abroad would be enough to absorb excess supply of good a. However, with a smaller appreciation of the terms of trade, the output produced in Country 1 looses in value, destroying the incentive to invest in the production capacity of Country 1 in the first place. Borrowing from abroad consequently cannot be an equilibrium solution.

In summary, it becomes clear that once the possibility of lower trade elasticities is accepted, a vast scope for policy intervention opens up in the regions where gains from the imposition of capital controls may be large. Moreover, in areas close to $\sigma(ToT)$, where the model exhibits explosive phenomena, the results call for policy intervention to anchor the system in a preferred
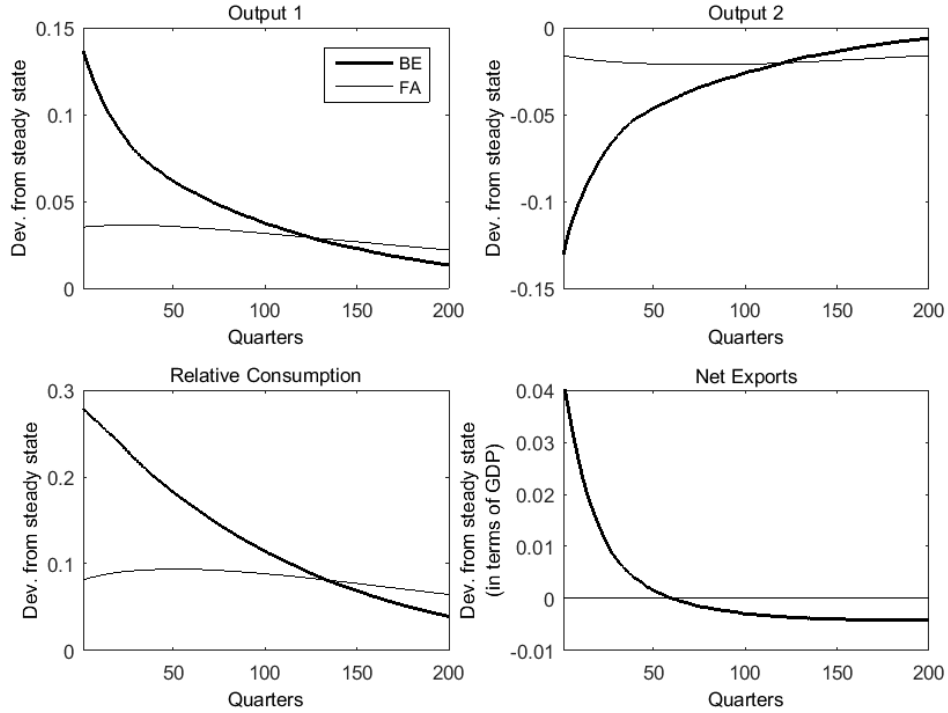
location either above or below $\sigma(ToT)$.

Figure 2.5: Impulse responses to a 1 standard deviation shock in Country 1's productivity. $\sigma = 0.1$, import share = 25%.

2.6.2 Continuous Capital Controls

So far, I have modelled capital controls as a complete shut-down in all asset trade. Although allowing for a clear-cut analysis, this might not be the best measure to look at. For one, capital controls come in many different forms in reality and only rarely as a complete shut-down in cross-border asset trade. Moreover, looking only at financial autarky might prevent insights on the possible magnitude in welfare gains from capital controls. As explained above, capital controls can cause a Pareto improvement, when asset trade under a bond economy is inefficiently high. It would consequently be intuitive to think that asset trade would optimally be dampened to an optimal level—but not cut off completely. Similar to a Laffer curve, the highest welfare gains would then lie somewhere between financial autarky and the bond economy. Interestingly, the results suggest otherwise.

Table 2.4 presents the welfare gains of imposing financial autarky for two representative calibrations of the trade elasticity, $\sigma = 0.5 > \sigma(ToT)$ and $\sigma = 0.1 < \sigma(ToT)$, given import shares of either 25% or 15%. Going across each row from left to right, the tax on bond returns increases from $\bar{\tau} = 0$ (BE) to $\bar{\tau} = 0.1$. As can be seen immediately, imposing financial autarky leads to a welfare gain *in every case*.[11] In other words, no intermediate solution Pareto-dominates

---

[11]Results show up as 0.000 due to rounding. All entries in Table 2.4 are strictly positive.

Welfare gains from imposing Financial Autarky							
$\sigma = 0.5$*	$\bar{\tau} = 0$ (BE)	$\bar{\tau} = 0.0001$	$\bar{\tau} = 0.001$	$\bar{\tau} = 0.01$	$\bar{\tau} = 0.1$	$\bar{\tau} = 0.5$	$\bar{\tau} = 1$
$is = 0.25$	0.060	0.043	0.022	0.008	0.001	0.000	0.000
$is = 0.15$	0.108	0.087	0.046	0.028	0.007	0.002	0.001
$\sigma = 0.1$**							
$is = 0.25$	3.714	3.328	1.978	0.574	0.101	0.026	0.014
$is = 0.15$	0.120	0.104	0.055	0.012	0.001	0.000	0.000

Table 2.4: Welfare gains calculated relative to bond economy with taxes as stated, as % of consumption; *unpruned solution, standard preferences; **pruned solution, Uzawa preferences. is = import share, σ = trade elasticity, $\bar{\tau}$ = constant component of tax rate.

a complete shut-down in asset trade. Taking a closer look, the higher the imposed tax, the smaller the welfare gain of imposing financial autarky. Increasing taxes leads to a step by step reduction in bond trade, converging towards financial autarky and thus the welfare level under financial autarky. This can also be seen in figure 2.6, where the imposition of taxes results in IRFs approaching financial autarky step by step but never breaking out of the band created by FA and BE.
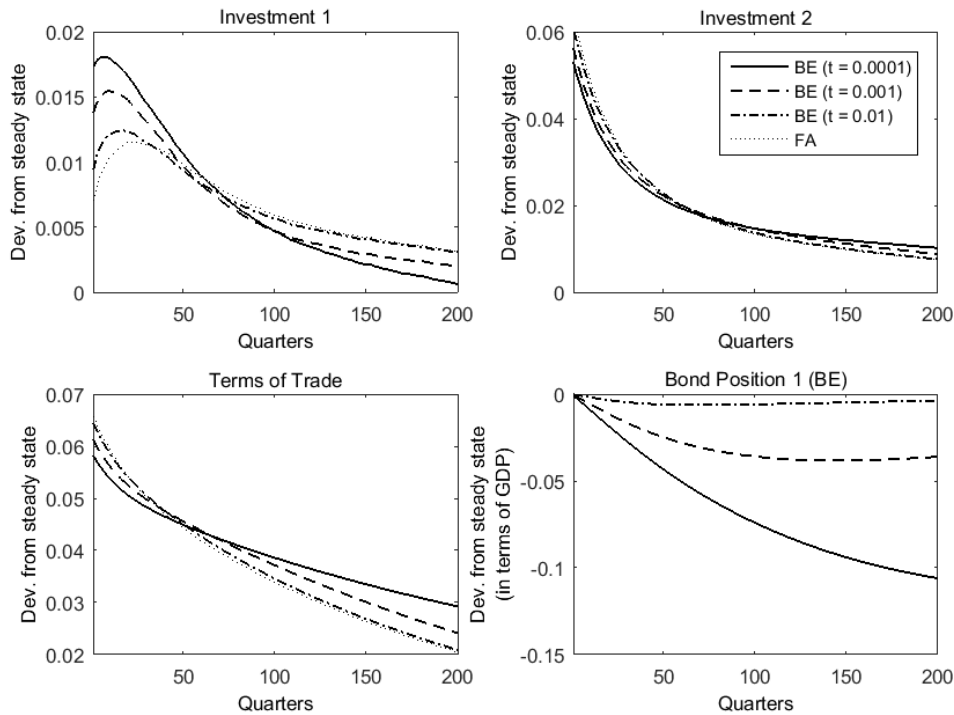

Figure 2.6: Impulse responses to a 1 standard deviation shock in Country 1's productivity. $\sigma = 0.5$, import share = 25%.

How can it be explained that even a small amount of asset trade is inferior to no asset trade? Firstly, by acknowledging that the tempting first conclusion of "no asset trade is better than some asset trade" needs an essential qualification. To the point, no asset trade is better than asset trade in the direction given by the BE equilibrium. As seen in Figure 2.6, at $is = 0.25$ and $\sigma = 0.5$, Country 1 enters a current account deficit by selling bonds to Country 2. Given the

above results, this flow in assets is not only inefficiently large, but actually going in the "wrong" direction.

This also explains, why a tax cannot improve upon financial autarky in the presented cases. Whilst taxes and tariffs can weaken or strengthen certain behavioural outcomes, they can not change their sign.

2.6.3   Asymmetric Countries

Up to this point, all calculations have been made with a calibration that produces two symmetrical countries. Whilst this is another step taken to simplify the exposition of underlying mechanisms, an empirically relevant analysis not only has to acknowledge the presence of country differences, but also check the robustness of its results in this respect. In what follows, I recalibrate the model twice. In the first instance, Country 1 takes on a size 1.5 times as large as Country 2. In the second instance, Country 1 is twice as large as Country 2.

Welfare gains from imposing financial autarky			
L = 1.5		$\sigma = 0.1^*$	$\sigma = 0.5^{**}$
$is = 0.25$	Country 1	14.13	-19.04
	Country 2	-9.08	24.49
$is = 0.15$	Country 1	13.94	-17.20
	Country 2	-10.22	18.57
L = 2		$\sigma = 0.1^*$	$\sigma = 0.5^{**}$
$is = 0.25$	Country 1	-3.60	-30.15
	Country 2	-0.83	45.51
$is = 0.15$	Country 1	24.70	-28.45
	Country 2	-15.50	31.72

Table 2.5: Welfare gains calculated relative to bond economy, as % of consumption, with Country 1 being 1.5 times (L=1.5) or 2 times (L=2) the size of Country 2; *pruned solution, Uzawa preferences; **unpruned solution, standard preferences.

As the two countries behave differently in an asymmetric setup, Table 2.5 presents the welfare gains of imposing financial autarky separately for each country. The top panel shows results with Country 1 being 1.5 times the size of Country 2, the bottom panel shows results for Country 1 being twice the size. For most calibrations, there is one country that clearly benefits from the imposition of capital controls. More importantly, however, there is *no case* in which both countries benefit from the introduction of capital controls. Put differently, there is no case where both countries would ex-ante agree to the introduction of capital controls. They never cause a Pareto improvement.

This result once again underlines that utmost caution is called for in dealing with capital controls and welfare comparisons. Whilst Pareto improvements could be substantial under a symmetric setup, the more realistic case of differing country sizes makes Pareto improvements seemingly much less plausible. It has to be added, that the results presented in Table 2.5 only

show a small selection of possible calibrations and there might be others where both countries experience an improvement through capital controls. In fact, taking a symmetric calibration where FA is Pareto superior to BE and infinitesimally changing country sizes still results in a Pareto improvement, showing that asymmetry and Pareto improvements are not mutually exclusive at all. Nonetheless, the results indicate that for significant differences in country size, improvements may be impossible to achieve.

This result extends not only to asymmetry in size, but also in business cycle characteristics. In particular, I simulated the interaction between a large developed country and a small emerging country by adjusting business cycle volatilities to empirically observed values. Whilst I specified the variance of innovations in the emerging country to $\sigma_\epsilon^2 = 0.02$, it was set to a significantly lower $\sigma_\epsilon^2 = 0.007$ for the developed country. Rerunning the above cases for this scenario produced results that were identical in every aspect except the exact magnitudes.

Finally, when looking at overall welfare for asymmetric countries, the result is mixed. Summing up welfare and weighting each country by its respective size results in an overall welfare gain in some cases, whereas others report a loss.

2.7 Conclusion

Using a two-country, two-good BKK framework, this paper has shed light on some of the welfare effects of capital controls. The analysis has produced three economic and one computational result.

Firstly, welfare effects are dominated by changing movements in the terms of trade across different trade elasticities. Not only do potential welfare gains from the imposition of capital controls grow as the terms of trade reaction to shocks grows, but also does the model behaviour fundamentally change for very low trade elasticities. All this points to potentially significant Pareto improvements through the imposition of capital controls.

Secondly, a moderate version of capital controls in the form of a tax on bond returns does not improve upon the welfare results under financial autarky. This indicates that in cases where capital controls can cause a Pareto improvement, market imperfections guide asset flows in the wrong direction. As a tax can only increase or decrease the volume of asset trade but not reverse flows, the best that can be done is a complete shut-down in bond trade.

Thirdly, with countries of significantly different size, introducing capital controls produces winners and losers. For all parametrisations solved in the course of the presented research, at least one country was clearly worse off from the imposition of financial autarky, making Pareto improvements impossible. This is a clear call for caution, showing that the robustness of Pareto-improving capital controls hinges crucially on assumptions about country size.

On a computational level, the paper has uncovered the powerful union of Uzawa preferences and pruning for solving problems related to explosiveness or lack of stationarity in macroeconomic models. Whilst the results were very promising, a host of questions was identified about the exact workings at play. Answering these questions took on a life of its own, with the aim of generalising the methods for a wider set of international macroeconomics applications.

On a general level, the results have shown that for certain parametrisations, countries would prefer to shut down bond trading ex-ante and thus increase cross-border risk sharing. In particular, the results have shown that the automatic insurance provided through terms of trade movements can be excessive in the face of pecuniary externalities. As agents overinvest following a productivity shock, the reaction of the terms of trade can be so large that the gain in productive efficiency is outweighed by a loss in risk sharing between the countries. When this is the case, imposing capital controls can result in a Pareto superior outcome.

Going beyond earlier research, this paper takes into account that the reaction of the terms of trade is discontinuous and non-monotonic over different intertemporal trade elasticities. Specifically, for low elasticities the impact response of the terms of trade to a productivity shock changes sign. Looking at the area of parametrisations where this is the case, the results indicate that welfare always improves when imposing financial autarky compared to a bond economy.

As shown in De Paoli and Lipinska (2013), models without productive capital do not feature the possibility for Pareto-improving capital controls. Going forward, it would be informative to further explore the exact role of (productive) capital and capital controls. In particular, (productive) capital was restricted in the above model as agents were only allowed to invest in their own country—even though the concept of capital controls did not touch productive capital. A valid point for further investigation would thus open up by looking at a model that allows agents to invest in and own foreign capital and how this affects asset prices.

Bibliography

Andreasen, M. M., Fernandez-Villaverde, J., and Rubio-Ramirez, J. F. (2013). The pruned state-space system for non-linear DSGE models: Theory and empirical applications. Working Paper 18983, National Bureau of Economic Research.

Backus, D. K., Kehoe, P. J., and Kydland, F. E. (1992). International real business cycles. *Journal of Political Economy*, pages 745–775.

Bhagwati, J. (1958). Immiserizing growth: a geometrical note. *The Review of Economic Studies*, pages 201–205.

Bianchi, J. (2011). Overborrowing and systemic externalities in the business cycle. *American Economic Review*, 101(7):3400–3426.

Bianchi, J. and Mendoza, E. G. (2013). Optimal time-consistent macroprudential policy. Working Paper 19704, National Bureau of Economic Research.

Bodenstein, M. (2011). Closing large open economy models. *Journal of International Economics*, (84):160–177.

Brunnermeier, M. K. and Sannikov, Y. (2015). International credit flows and pecuniary externalities. *American Economic Journal: Macroeconomics*, 7(1):297–338.

Caballero, R. J. and Krishnamurthy, A. (2004). Smoothing sudden stops. *Journal of Economic Theory*, 119(1):104 – 127.

Cole, H. L. and Obstfeld, M. (1991). Commodity trade and international risk sharing: How much do financial market matter? *Journal of Monetary Economics*, 28:3–24.

Corsetti, G., Dedola, L., and Leduc, S. (2008). International risk sharing and the transmission of productivity shocks. *The Review of Economic Studies*, 75(2):443–473.

Costinot, A., Lorenzoni, G., and Werning, I. (2014). A theory of capital controls as dynamic terms-of-trade manipulation. *Journal of Political Economy*, 122(1):77 – 128.

De Paoli, B. and Lipinska, A. (2013). Capital controls: A normative analysis. Staff Report 600, Federal Reserve Bank of New York.

Den Haan, W. J. and De Wind, J. (2012). Nonlinear and stable perturbation-based approximations. *Journal of Economic Dynamics and Control*, 36(10):1477 – 1497.

Eichengreen, B. and Rose, A. (2014). Capital controls in the 21st century. *Journal of International Money and Finance*, 48(PA):1–16.

Farhi, E. and Werning, I. (2014a). Dilemma not trilemma? Capital controls and exchange rates with volatile capital flows. *IMF Economic Review*, 62:569–605.

Farhi, E. and Werning, I. (2014b). Fiscal unions. Working Paper 198816, Harvard University OpenScholar.

Forbes, K., Fratzscher, M., Kostka, T., and Straub, R. (2012). Bubble thy neighbor: Portfolio effects and externalities from capital controls. Working Paper 18052, National Bureau of Economic Research.

Heathcote, J. and Perri, F. (2013). The international diversification puzzle is not as bad as you think. *Journal of Political Economy*, 121(6):1108 – 1159.

Heathcote, J. and Perri, F. (2014). On the desirability of capital controls. Working paper, Federal Reserve Bank of Minneapolis.

Helpman, E. and Razin, A. (1978). A theory of international trade under uncertainty. MPRA Paper 22112, University Library of Munich, Germany.

Hooper, P., Johnson, K., and Marquez, J. (2000). Trade elasticities for the G-7 countries. Princeton Studies in International Economics 87, Princeton University.

Judd, K. L. (1998). *Numerical Methods in Economics*. MIT Press, Cambridge, MA.

Keynes, J. M. (1929). The German transfer problem. *The Economic Journal*, 39:1–7.

Kim, J., Kim, S., Schaumburg, E., and Sims, C. A. (2008). Calculating and using second-order accurate solutions of discrete time dynamic equilibrium models. *Journal of Economic Dynamics and Control*, 32(11):3397 – 3414.

Korinek, A. (2011). The new economics of capital controls imposed for prudential reasons. IMF Working Papers 11/298, International Monetary Fund.

Lambert, F., Ramos-Tallada, J., and Rebillard, C. (2011). Capital controls and spillover effects: evidence from Latin-American countries. Working papers 357, Banque de France.

Lewis, K. K. (1999). Trying to explain home bias in equities and consumption. *Journal of Economic Literature*, 37(2):571–608.

Lombardo, G. and Uhlig, H. (2014). A theory of pruning. ECB Working Paper 1696.

Magud, N. E., Reinhart, C. M., and Rogoff, K. S. (2011). Capital controls: Myth and reality - a portfolio balance approach. Working Paper 16805, National Bureau of Economic Research.

Martin, A. and Taddei, F. (2013). International capital flows and credit market imperfections: A tale of two frictions. *Journal of International Economics*, 89(2):441 – 452.

Neumeyer, P. A. and Perri, F. (2005). Business cycles in emerging economies: the role of interest rates. *Journal of Monetary Economics*, 52(2):345–380.

Obstfeld, M. and Rogoff, K. (1996). *Foundations of International Macroeconomics*. The MIT Press, 1st edition.

Ohlin, B. (1929). Mr. Keynes' views on the transfer problem. II. A rejoinder from Professor Ohlin. *The Economic Journal*, 39:400–4.

Ostry, J. D., Ghosh, A. R., Habermeier, K., Chamon, M., Qureshi, M. S., and Reinhardt, D. B. (2010). Capital inflows: The role of controls. IMF Staff Position Note SPN/10/04, IMF.

Ostry, J. D., Ghosh, A. R., and Korinek, A. (2012). Multilateral aspects of managing the capital account. IMF Staff Discussion Note SDN/12/10, IMF.

Schmitt-Grohe, S. and Uribe, M. (2003). Closing small open economy models. *Journal of International Economics*, 61(1):163–185.

Schmitt-Grohe, S. and Uribe, M. (2007). Optimal simple and implementable monetary and fiscal rules. *Journal of Monetary Economics*, 54(6):1702–1725.

Taylor, J. (1993). *Macroeconomic Policy in a World Economy: From Economic Design to Practical Operation*. Norton, New York, NY.

Uzawa, H. (1968). Time preference, the consumption function and optimum asset holdings. In Wolfe, J. N., editor, *Value, Capital and Growth: Papers in Honor of Sir John Hicks*, pages 485–504. The University of Edinburgh Press.

World Bank (2015). Database: Imports of goods and services. `http://data.worldbank.org/indicator/NE.IMP.GNFS.ZS?order=wbapi_data_value_2014+wbapi_data_value+wbapi_data_value-last&sort=asc`. Retrieved on 22.06.2015.

A Equilibrium Conditions

Period utility function

$$U(c_i, n_i) = \log c_{it} - \phi \frac{n_{it}^2}{2}.$$ (33)

Complete Markets

Intertemporal Euler equations for capital

$$\frac{1}{c_1(s^t)} = \beta \frac{1}{c_1(s^{t+1})} \left(r_1(s^{t+1}) q_1^a(s^{t+1}) + 1 - \delta \right)$$ (34)

$$\frac{1}{c_2(s^t)} = \beta \frac{1}{c_2(s^{t+1})} \left(r_2(s^{t+1}) q_2^b(s^{t+1}) + 1 - \delta \right)$$ (35)

Intratemporal Euler equations for consumption and leisure

$$\phi n_1(s^t) = \frac{w_1(s^t) q_1^a(s^t)}{c_1(s^t)}$$ (36)

$$\phi n_2(s^t) = \frac{w_2(s^t) q_2^b(s^t)}{c_2(s^t)}$$ (37)

International risk sharing conditions for a and b

$$\frac{q_1^a(s^t)}{c_1(s^t)} = \frac{q_2^a(s^t)}{c_2(s^t)}$$ (38)

$$\frac{q_1^b(s^t)}{c_1(s^t)} = \frac{q_2^b(s^t)}{c_2(s^t)}$$ (39)

Bond Economy

Intertemporal Euler equations for bonds

$$\frac{1}{c_1(s^t)} P(s^t) = \beta \frac{1}{c_1(s^{t+1})} \left(\frac{1}{2} + \frac{1}{2} rx(s^{t+1}) \right)$$ (40)

$$\frac{1}{c_2(s^t)} \frac{P(s^t)}{rx(s^t)} = \beta \frac{1}{c_2(s^{t+1})} \left(\frac{1}{2} \frac{1}{rx(s^{t+1})} + \frac{1}{2} \right)$$ (41)

Including taxes

$$\frac{1}{c_1(s^t)} P(s^t) = \beta \frac{1}{c_1(s^{t+1})} \left(\frac{1}{2} + \frac{1}{2} rx(s^{t+1}) \right) (1 + \tau_1(s^t))$$ (42)

$$\frac{1}{c_2(s^t)} \frac{P(s^t)}{rx(s^t)} = \beta \frac{1}{c_2(s^{t+1})} \left(\frac{1}{2} \frac{1}{rx(s^{t+1})} + \frac{1}{2} \right) (1 + \tau_2(s^t))$$ (43)

B Finding $\sigma(ToT)$

Log-linearising the set of model equations, one can show that the deviation of the terms of trade from steady state under financial autarky are given by

$$\hat{p} = \frac{\hat{y}_1 - \hat{y}_2}{1 + 2s(\sigma - 1)} \tag{44}$$

where $\hat{x}$ denotes the percentage deviation of variable x from its steady-state value and s is defined as the share of locally produced intermediate goods in the final goods production.

It follows that, given a positive shock to output in Country 1, the function changes sign at

$$\sigma(ToT) = 1 - \frac{1}{2s}. \tag{45}$$

For the general case of a BE, movements in the terms of trade are given by

$$\hat{p} = \frac{1}{1 + 2s(\sigma - 1)}\left(\frac{2s - 1}{1 - s} \times nx + \hat{y}_1 - \hat{y}_2\right). \tag{46}$$

In steady state, the relationship between s and the home-bias parameter *omega* is given by

$$s = \frac{\omega^\sigma}{\omega^\sigma + (1 - \omega)^\sigma}. \tag{47}$$

Chapter 3

Bridging Fundamental and Market Sentiment Approaches in Emerging Market Currency Valuation

3.1 Introduction

Predicting exchange rate movements is notoriously difficult. Famously, Meese and Rogoff (1983a) established the result that a random walk outperforms forecasts based on economic models of exchange rates. Ever since, contributions to the literature have proposed models that could improve on the random walk when tested out of sample. As the more recent overview by Rogoff (2001) asserts, the original result remains largely untouched. A notable exception comes in the form of very long-horizon forecasts out to three to four years, which can do significantly better than the random walk, as shown prominently in Mark (1995), amongst others.

This leaves the question of how to progress on the short and medium end of the forecasting horizon. As pointed out in detail in Section 3.2, countless contributions to the literature have taken steps in this direction. Yet despite improving our understanding of the behaviour of exchange rates, the proposed models seem to suffer from three marked deficiencies when applied out of sample. Firstly, many models are highly sensitive to the sample period and produce unsatisfying results once earlier or later periods are taken into account. Secondly, many applied exchange rate forecasts seem to be based on an arbitrary selection of explanatory variables, justified by the finding that "they work". Anecdotal evidence suggests a particular prevalence of this type of model construction in the FX research departments of investment banks, where different currencies are often forecast using completely different sets of explanatory variables such as commodities, equities, CDS spreads, yield curves and others. While each of these variables may justifiably form part of a forecast, using a different set of variables for each currency reduces

the generality and appeal of an approach. Lastly, a number of technically more sophisticated forecasting approaches, e.g. models built on the factor model approach of Engel et al. (2012), have a tendency to obscure the exact composition of a given forecast. Put differently, when a forecast is generated not only to mechanically invest, but also to provide information on the nature and behaviour of a currency, then having a black-box mechanism can obscure the insights a researcher wishes to gain.

This paper assesses whether we can correctly forecast exchange rate excess returns over the short-horizon by combining the literature on fundamental exchange rate forecasting with non-fundamental and market-sentiment oriented approaches. In doing so, we aim to address the shortcomings mentioned above. Concretely, we construct a model that employs a uniform methodology across a sample of 20 emerging market currencies. The model not only applies the same approach across all countries, but also leaves room for the data to speak for itself when it comes to deciding how important certain building blocks of the model are and in which manner they influence the forecast. This allows the model to adapt to changing market circumstances while at the same time minimising arbitrary choices of the forecaster. Concerning the exact econometric specification, we decide to choose data that allows the use of the simple yet powerful panel OLS. In doing so, the output of the model remains straightforward to interpret, giving the researcher not only a signal on which to base a trade, but also a host of information on the state of the currency and its dynamics.

The main innovation of the paper lies in bridging existing strands on currency forecasting. Based on the observation in Menkhoff and Taylor (2007) that it is necessary to take into account both fundamental and non-fundamental aspects when making informed decisions about currencies, we build a model that includes both from the start. A first step in this direction has been taken by Jorda and Taylor (2009), but the remarkable majority of models still treat fundamental and non-fundamental approaches as separate. While it can be useful for an investor to have different sets of forecasts based on different approaches, this poses the problem of how to weight them when making an investment decision. Our model overcomes this problem by producing only one forecast based on a combination of the different approaches.

Concretely, we include four main building blocks. Firstly, a measure of real exchange rate strength captures fundamental value in the model and gives it traction in the long run. Secondly, the model captures shorter-term dynamics by including the nominal interest rate differential—of particular importance due to the prevalence of currency carry trades. Thirdly, the model falls back on contributions about the interplay of equity markets and currencies, as well as technical analysis, by including local equity returns. This measure in itself bridges technical concepts such as "momentum" and "reversal" with the fundamental buying and selling of foreign currencies to invest in foreign equity. Lastly, our model includes a measure of economic surprises to cover

local market sentiment in the very short-run.

While remaining parsimonious, the model thus bridges a variety of modelling approaches and builds on insights from various strands of the literature developed over past decades.

The rest of the paper is structured as follows. Section 3.2 gives an overview of the existing literature. Section 3.3 describes a theoretical background for modelling currency risk premia, while Section 3.4 presents the model itself. Additional results and robustness are addressed in Section 3.5. Section 3.6 concludes.

3.2 Related Literature

It is safe to assume that efforts were made to understand currency value and exchange rates, in one form or another, since the inception of internationally used currencies. That said, most current research builds on a number of seminal contributions from the second half of the 20th and the beginning of the 21st century. In fact, the field has seen an enormous proliferation of research in the last decades, creating a diverse landscape of models and methodologies. For the purpose of this paper, we decide to give an overview only of those branches of the literature which are of immediate relevance to the presented research,[1] namely, fundamental exchange rate determination and more recent technical and sentiment-driven approaches to exchange rates.

For fundamental exchange rate research, the nihilistic paper by Meese and Rogoff (1983a) provided an era-defining contribution, stating that exchange rate forecasts could not beat a random walk prediction. The fundamental or "value" models in question would typically try to predict future exchange rate movements on the basis of economic and monetary fundamentals, such as current account imbalances or real exchange rate misalignments. Contributions in the following decades largely confirmed the result of Meese and Rogoff, with one significant exception. Mark (1995) added the crucial qualification that models seem to outperform the random walk at very long horizons of more than three years. Despite this qualification, the general mood in the field remained depressed, as expressed, for example, by Sarno and Taylor (2002): "[E]mpirical work on exchange rates still has not produced models that are sufficiently statistically satisfactory to be considered reliable and robust". More recently, this sentiment was directly addressed by Engel et al. (2008) with the telling title "Exchange Rate Models Are Not As Bad As You Think". Rather than proposing a distinct new model, however, the authors argue that "beating a random walk" is too strong a criterion in evaluating forecasts and show that out-of-sample forecasting power can be increased with panel estimation as well as by focusing on long-horizon forecasts. This leaves, as Rogoff (2008) puts it, "the big challenge: one month to one year horizons".

---

[1]For a broader overview on exchange rate research, the interested reader is guided to James et al. (2012).

As an alternative to fundamental or value models, the long-standing strand of the literature termed "technical analysis" uses a currency's own past movements to predict future returns in a rule-based manner. As summarised in Pring (1991), technical analysis attempts to profit from systematic changes in the psychology of the market. In their review of technical analysis in currency markets, Menkhoff and Taylor (2007) document that technical analysis is widespread and continually used, but that it remains unclear whether the realisation of profits from technical analysis stems from incurring large risks which might not be fully compensated for.

A third strand of the literature of relevance to our analysis focuses on the widespread application of "carry trade" strategies and their implications on exchange rates. Based on the empirical observation that high-yielding currencies often do not depreciate as predicted by uncovered interest parity theory, carry strategies generate returns by borrowing in low interest rate currencies and investing in high interest rate currencies. Burnside et al. (2011a) argue that so-called Peso problems, i.e. low probability events that do not occur in the sample, explain these returns, even when negative payoffs during Peso periods are not very large. More recently, Koijen et al. (2016) show that although other known predictors of returns do not explain carry, the latter captures a number of known predictors from different asset classes, making it useful to the prediction of future prices.

Realising the potential of combining some of these strands, Jorda and Taylor (2009) take a first step in creating a model that combines the carry trade with a fundamental approach. Specifically the authors condition their model's carry trade recommendations on fundamental equilibrium exchange rates (FEERs), achieving better returns as measured by Sharpe ratios and skewness when forecasting currencies for nine developed countries out of sample.

Finally, a number of contributions have realised the potential to increase model performance by building currency portfolios instead of forecasting bilateral exchange rates independently. Examples include Lustig and Verdelhan (2007), Lustig et al. (2011), Menkhoff et al. (2011) and Menkhoff et al. (2016).

3.3 Theoretical Background for Modelling Currency Risk Premia

To set the stage for our empirical analysis, this section introduces the theoretical concept of risk premia, or expected excess returns.

If S_t is the nominal exchange rate, defined in units of home currency per unit of foreign currency, F_t the one-period forward rate, and i_{t+1} and i_{t+1}^* the nominal interest rates at home

and abroad, then the covered interest parity (CIP) condition ensures that

$$1 + i_{t+1} = (1 + i_{t+1}^*)\frac{F_t}{S_t}.$$ (3.1)

Moreover, the uncovered interest parity (UIP) condition formalises the idea that the return earned on investments in different currencies should, in expectation, be identical. Algebraically,

$$1 + i_{t+1} = (1 + i_{t+1}^*)E_t\Big[\frac{S_{t+1}}{S_t}\Big],$$ (3.2)

where i and i^* denote the nominal interest rates at home and abroad, respectively, and S is the nominal exchange rate defined as units of home currency per unit of foreign currency (see Obstfeld and Rogoff (1996) for a textbook analysis of CIP and UIP). The UIP condition implies that any differences in nominal interest rates are expected to be balanced out by an appreciation or depreciation in the bilateral exchange rate. The observation that realised exchange rate movements repeatedly fail to do so is often misinterpreted as a failure of UIP. Note, however, that the condition describes an ex ante, not an ex post relationship. As described in Jorda and Taylor (2009), evidence points to ex ante exchange rate expectations not being far out of line with interest differentials. The problem rather seems to stem from systematically wrong expectations themselves. Ex post, these drive a wedge into the realised returns of investments in different currencies.

Defining the risk premium as

$$rp_t = \frac{1 + i_{t+1}^*}{1 + i_{t+1}}E_t\Big[\frac{S_{t+1}}{S_t}\Big],$$ (3.3)

any deviation of rp_t from 1 would indicate a violation of UIP. Ex post, the realised excess return can be defined as

$$rx_t = \frac{1 + i_{t+1}^*}{1 + i_{t+1}}\frac{S_{t+1}}{S_t} = \frac{S_{t+1}}{F_t},$$ (3.4)

where the last equality directly follows from the CIP condition 3.1.

Using the definition for the currency risk premium allows us to distinguish between its two main components: the interest differential and expected changes in the nominal exchange rate S_t. We take a closer look at each in turn.

There is a diverse literature looking at the influence of interest rate differentials on the determination of exchange rates, as emphasised by Engel and West (2010). Our approach here is mostly definitional. It is evident from the definition of rx that, if the nominal exchange rate were to stay completely unchanged from one period to the next, the interest rate differential would constitute the entirety of the risk premium. In the literature, the concept of the income

earned from holding an asset if its price stays the same is known as "carry". Koijen et al. (2016) note that carry is special in so far as it is a model-free characteristic that is observable ex ante. This property makes it an indispensable component of any risk premium forecast. Were we to assume, as is regularly done, that exchange rates follow a random walk, i.e. $E_t(S_{t+1} - S_t) = 0$ (see e.g. Meese and Singleton (1982) for an early example), then our forecast of the risk premium would indeed be identical to the interest rate differential. How central the idea of carry is to applied currency management and trading can also be seen in the prevalence of the "carry trade", a strategy solely focussed on capturing returns from interest rate differentials that have not been neutralised by counteracting nominal depreciations.

In practice, finding reliable data on interest rates at a regular frequency can pose a challenge, especially for emerging market countries. To calculate the interest rate differential we can, however, circumvent this problem. From the CIP condition 3.1 it follows that the interest rate differential can be expressed as the difference between forward and spot prices, both of which are reliably and uniformly available in the market (see more in Section 3.4):

$$\frac{1 + i^*_{t+1}}{1 + i_{t+1}} = \frac{S_t}{F_t}. \tag{3.5}$$

The second main component of the risk premium is given by the change in the nominal exchange rate $\frac{S_{t+1}}{S_t}$. The literature broadly distinguishes between two approaches to modelling changes in the exchange rate. The first tries to look at the underlying value of a currency and consequently focuses on the real exchange rate and its "fair value". Recent contributions have shown that the real exchange rate is stationary, implying there is some long-run value to which it eventually converges. However, this reversal towards the long-run average is very slow with estimates for the half-life of real exchange rate deviations ranging from 3 to 5 years (Rabanal and Rubio-Ramirez, 2015). For this reason, the fair value approach is predominantly chosen to predict exchange rate movements over a long-horizon. Typical examples in practice (such as Trivedi and Ozerov (2016)) gain traction 4 to 8 quarters out, and become most powerful at 12 to 16 quarters ahead.

In their most basic form, fundamental models base forecasts on deviations from purchasing power parity (PPP) theory, which states that real exchange rates should equal 1, or at least have a tendency to revert back to 1 when that long-run ratio is disturbed for some reason (Obstfeld and Rogoff, 1996). As mentioned above, however, these deviations can be very persistent. Consequently, more differentiated models try to identify which deviations are justified, e.g. by fundamentals such as productivity growth, and which are not. Examples include models of fundamental equilibrium exchange rates (FEER) that pin down the level of a currency that is consistent with achieving both a sustainable current account deficit and output close to potential,

or so-called behavioural equilibrium exchange rates (BEER) that use fundamentals directly to estimate the equilibrium level of exchange rates (Clark and MacDonald, 1999).[2]

The second school of thought in forecasting currency movements builds on the observation that fundamentals have close to no predictive power over short time horizons, a result first popularised by Meese and Rogoff (1983a,b). Instead, short-term movements seem to be driven by investor sentiment and various other manifestations of behavioural anomalies, often in conflict with the efficient market hypothesis. Amongst the most widely followed strategies are "carry", "momentum" and "reversal". With an extensive literature covering these strategies (and others) this paper will not analyse them in depth, but rather give a short overview and focus on the concepts underlying the usefulness of the strategies.

Carry strategies are based on the empirical observation that nominal exchange rates often do not adjust for interest rate differentials across borders. Borrowing in low-yielding currencies and investing in currencies with high yields can then leave investors with an excess return. Burnside et al. (2011a) and Burnside et al. (2011b), amongst others, document the success of this strategy, but also warn of the property of carry trades to incur large and rapid losses during times off financial market turmoil. As Barroso and Santa-Clara (2015) point out, "high-yielding currencies are known to 'go up by the stairs and down by the elevator', implying that the carry trade has substantial crash risk". Nonetheless, carry remains a popular strategy amongst institutional investors and private individuals alike.[3] As mentioned in Section 3.2, Koijen et al. (2016) show that carry also captures several known return predictors from different asset classes, implying its usefulness for forecasting.

Momentum strategies are based on empirical observations that periods of currency moves in one direction are sometimes systematically followed by further moves in the same direction. The closely related reversal strategies, on the other hand, profit from FX movements that change direction after a certain period of time. The former is rationalised by a crowding-in behaviour, where investors seek out currencies that were successful in the past and buy into them leading to further appreciations. Conversely, the latter strategy can be rationalised with a "cashing-in" behaviour that leads investors to sell currencies after a period of appreciation to secure realised profits (see Pring (1991), Neely and Weller (2012), Menkhoff et al. (2011), Burnside et al. (2011b) and Asness et al. (2013)). Whilst a number of empirical studies come up with successful versions of momentum and reversal strategies, they have significant shortcomings. Most importantly, returns from momentum and reversal strategies display a high degree of sensitivity

---

[2]To find a way through the proliferation of acronyms stemming from various approaches—CHEER, ITMEER, PEER, DEER, FEER and BEER, to name just a few—the interested reader is guided to Driver and Westaway (2004).

[3]For an account of the surprising extent to which the carry trade became popular amongst Japanese housewives, as well as the personal tragedies following episodes during which carry incurred large losses, see Jorda and Taylor (2009).

to chosen sample periods and currency pairs. Even though many versions of these strategies produce positive excess returns over a certain number of years, they then regularly fall victim to changes in market dynamics, often wiping out all previous gains (see Menkhoff and Taylor (2007) for an analysis of the question in how far profits are rationalised by large and uncompensated risk). One conclusion to be drawn is that these strategies might not be beneficial on their own when forecasting exchange rates over prolonged periods of time, but may add useful information to a more sophisticated model, at least during certain periods. Moreover, to be convincing not only ex post, but also ex ante, a methodologically appealing setup would apply the strategy uniformly to all currency pairs and retain a certain level of agnosticism about the exact time periods and directions involved.

Among the many other forms of behavioural investment strategies, such as looking at investor flows as in Froot and Ramadorai (2005), of particular importance is the effect of market sentiment itself. Traditionally, there are currencies which profit from a so-called "risk-on" environment, such as most emerging market currencies. Other currencies, in particular those known as "safe havens", such as the Swiss franc or the Japanese yen, strengthen in times of uncertainty—so-called "risk-off" episodes. This leads many forecasters to introduce dummies that push forecasts in one way or another during times of high or low volatility, depending on the "character" of the currency. This approach is straightforward for currencies such as the Swiss franc, where a broad consensus exists as to its safe haven status. Such classifications become more difficult for many other currencies, however, particularly in emerging markets. Intuitively, there are two reasons to explain this. Firstly, whenever the level of political stability changes, so does the extent to which a currency is exposed to market turmoil and worldwide volatility. As political stability increases and decreases at a much higher frequency in emerging markets than in developed markets, categorisations, even when correct, are less long-lived. Secondly, emerging market currencies react very differently to different kinds of shocks and volatility. Whereas general USD strength affects emerging markets across the board, the effects are much more differentiated following shocks to the oil price and other commodities, to armed conflict or to demand shocks in big economic powers such as the US or China. Consequently, it is again paramount not to use a certain model setup that might provide good results in sample, but does not take into account that dynamics can change quickly, leading to bad forecasts later on.

In what follows, the strategy is to put the fundamental approach at the core of modelling risk premia. Doing so provides a foundation for prediction accuracy over a long horizon, independent of day-to-day market turbulence and changes in sentiment. That said, the paper poses itself the veritable challenge of forecasting risk premia over the short time span of 1 month. This is a

novelty in the literature on fundamental exchange rate forecasting, which typically sees success over horizons of 3 to 5 years. To be able to push the boundary of the investment horizon to such an extent, the paper explores new territory by combining the fundamental approach with the short-term investment strategies presented above. In particular, the aim is to provide a framework which is not only completely uniform in how it applies to the chosen currency pairs, but also remains agnostic about the way in which market sentiment and behavioural aspects affect the forecast. In this sense, the data is allowed to speak for itself and parameters will have the opportunity to adapt to changing market circumstances without being arbitrarily restricted to an approach that "works" in sample.

3.4   The Empirical Model

This section explains in detail which indicators are chosen to form the building blocks of the model, following the forecasting schools presented above. In particular, we will address the shortcomings of these schools and how to overcome them.

Fundamental Value: The Real Exchange Rate

Real exchange rate strength builds the fundamental core of our approach, in line with the established literature on real exchange rate modelling. As mentioned in the section above, one of the main difficulties lies in identifying real exchange rate deviations that are not justified by other fundamentals such as productivity growth and hence can be expected to revert in the near future. To avoid having to impose a way of estimating equilibrium real exchange rate levels across countries, we choose instead to focus on deviations of a country's real exchange rate from its own history. Amongst others, Engel (2016) finds strong evidence of mean reversion in the real exchange rate, implying that current deviations from the historical mean contain significant information about future movements, even when leaving aside cross-country comparisons. As RER deviations correct only very slowly—estimates range from 3 to 5 years (Rabanal and Rubio-Ramirez, 2015)—we use deviations from the ten-year average. Specifically, we calculate the 10-year z-score of the CPI-based broad real effective exchange rate constructed by JP Morgan. Using this measure makes observations from different countries comparable as the z-score corrects deviations for the general level of volatility in a country's time series.

Carry: The Interest Rate Differential

Using equation 3.4, we substitute the interest rate differential with the forward premium S_t/F_t. Data for both spot and forward rates are readily available for all countries in our sample. For forward rates, we choose 1-month non-deliverable contracts, corresponding to our forecast

horizon.

In choosing an intuitively convincing measure of the interest rate differential, the question arises whether the real interest rate differential might contain more information than the nominal. We make our choice on the basis of existing analysis which finds a higher information content in nominal compared to real interest rate differentials when it comes to predicting excess returns (see The Economist (2013) and references therein).

Momentum: Equity Returns

Momentum and reversal strategies traditionally focus on using the past performance of a currency to predict its future performance. In line with the more fundamental-oriented approach of this paper, however, we move away from this purely self-deterministic approach and use past performance of local equity markets as an explanatory variable for risk premia. This idea is based on Cenedese et al. (2015), who suggest that the correlation between international equity returns and currency returns could be positive due to return-chasing by investors. Specifically, they mention a large literature showing that investors often increase their holdings in markets that have recently outperformed, a phenomenon termed "trend chasing". That said, the authors also mention conflicting evidence pointing towards a negative correlation following the sell-off of recently successful stocks to repatriate returns. We consequently introduce equity returns in an agnostic way, allowing for either a positive or negative correlation in our forecasts. Specifically, we use the 1-month percentage change in the local MSCI equity index for each currency.

Investor Sentiment: The Citi Economic Surprise Index

The final building block of our forecast is given by the Citi Economic Surprise Index (CESI). Exchange rates are known to react almost instantaneously to announcements of macroeconomic news, such as new data (see Engel et al. (2008) and references therein). As such, most fundamental information is priced into exchange rates straight away. Nonetheless, positive or negative economic surprises can significantly affect investor sentiment with regards to a certain country or currency. Current sentiment, in turn, partly determines an investor's willingness to take on risk. This effect is particularly dominant in emerging markets, where large capital flows have been shown to sometimes follow sentiment, rather than fundamentals, not least during times of currency crises.[4] For exchange rate forecasting, this observation implies that economic surprises potentially contain systematic information about future moves in the exchange rate and risk premia. The CESI is by construction centred around zero, allowing us to use the index level directly for our estimation. Specifically, we use the CESI level from the last day of each month.

---

[4]See Flood et al. (2012) for a summary on the dynamics of currency crises.

Before moving on to the next section it is worth stressing again that despite including a large number of emerging market currencies, our approach is uniform across all countries. For no currency pair do we modify the number of regressors or how the parameters filter through to the forecast. This is an important property in the field of exchange rate forecasting, which, especially in market application, is ridden with traditions of arbitrarily including different regressors for different currency pairs until the model "works".

3.4.1 Data

Our sample contains monthly data for 20 emerging market economies: Brazil, Chile, Colombia, Czech Republic, Hungary, India, Indonesia, Israel, Malaysia, Mexico, Peru, Philippines, Poland, Russia, Singapore, South Africa, South Korea, Taiwan, Thailand and Turkey. All exchange rates in the sample are included as bilateral against the US dollar. The chosen currencies correspond to the most widely traded currencies in emerging markets as reported by the Bank for International Settlements (2016), with the exception of the Chinese yuan (renminbi), Hong Kong dollar and Saudi riyal, which we excluded due to their closely managed or pegged nature. The sample starts at the point in time from which data for the chosen variables exist for all countries, January 2005, and ends in December 2016.

Bilateral spot exchange rates, 1-month non-deliverable forwards, JP Morgan CPI-based broad real effective exchange rates and MSCI local equity indices are obtained from Bloomberg, capturing values at New York close on the last trading day of each month. Data for the month-end level of Citi Economic Surprise Indices are obtained from Citi.

3.4.2 Estimation

The estimation approach follows three main ideas. Firstly, we want the data to speak for itself such that sign and size of the coefficients can change over time. We achieve this through an extending window regression where every new observation is used to re-estimate coefficients for the following period. Secondly, we try to avoid a "black-box" mechanism at all cost. Looking for the most intuitive and easy to interpret estimation procedure available, we rely on ordinary least squares (OLS). As specified in the section above, all our data either naturally fulfil requirements that make OLS a valid approach, or they have been transformed to do so—e.g. by taking the z-score to account for the variation in standard deviations across countries. Thirdly, we aim to estimate parameters robustly despite the short time series available for emerging market data. We find that a straightforward way of doing so is by pooling our data and thus significantly increasing the number of observations. Mark and Sul (2012) show that, in this context, pooling dominates time-series regression when the heterogeneity in model parameters is small. By restricting our analysis to emerging market currencies that exhibit very similar properties, we

believe this condition to be met.

We estimate risk premia as

$$rp_{i,t+1} = \alpha + \beta_1 rer_{i,t} + \beta_2 ird_{i,t} + \beta_3 emom_{i,t} + \beta_4 cesi_{i,t} + u_{i,t}, \qquad (3.6)$$

where rp is the risk premium, rer is the real exchange rate z-score, ird is the interest rate differential, $emom$ is equity momentum, $cesi$ is the Citi Economic Surprise Index and u is the disturbance.

Table 3.1 shows the regression output from the in-sample period Jan 2005 to Dec 2005, including 217 observations. Despite the very short sample period, it already becomes clear that the interest rate differential is crucial to forecasting risk premia. The coefficient is positive and significant at the 1% confidence level. As expected, real exchange rate strength has a negative effect, albeit not significant at this stage. Equity returns have a negative coefficient, supporting the concept of 'reversal' while positive economic surprises have a positive coefficient. That said, again neither of the two is significant in this sample. At 8.9%, the R-squared is unusually high for a predictive equation of this type, potentially indicating that the interest rate differential explained an abnormally large amount of risk premia over 2005.

Table 3.1: In-sample estimation of risk premia.

Dependent Variable: risk premium
Method: Panel Least Squares
Sample (adjusted): 2005M01 2005M12
Periods included: 11
Cross-sections included: 20
Total panel(unbalanced) observations: 217

Variable	Coefficient	Std. Error	t-Statistic	Prob.
c	-0.15465	0.18129	-0.85305	0.3946
rer	-0.18437	0.140349	-1.31363	0.1904
ird	1.630068***	0.380886	4.279679	0.0002
emom	-0.02804	0.023441	-1.19638	0.2329
cesi	0.002654	0.002089	1.270495	0.2053
R-squared	0.089478	Adjusted R-squared		0.072298

3.4.3 Portfolio Strategy

In the next section we implement portfolio strategies based on the framework presented above.

For the construction of portfolios we follow the literature, in particular Menkhoff et al. (2016). We construct linear, rank and high-minus-low portfolios, using the forecasts for risk premia as signals.

For our benchmark linear portfolio, weights are assigned according to

$$w_{j,t+1} = c_t(rp_{j,t+1} - \overline{rp}_{t+1}), \tag{3.7}$$

where rp is the estimated risk premium and $\overline{rp}_{t+1} = \frac{1}{n}\sum_{j=1}^{n} rp_{j,t+1}$ is the cross-sectional average of estimated risk premia across countries n. The factor c_t is constructed such that the sum of all positive portfolio weights equals one and the sum of all negative portfolio weights equals one. In other words, the portfolio prescribes to borrow and invest one unit. This not only ensures comparability of portfolios across different set-ups, but also makes sure that the amount invested remains constant over time, independent of accumulated profits or losses. Algebraically,

$$c_t = \frac{1}{\sum_j |rp_{j,t+1} - \overline{rp}_{t+1}|}. \tag{3.8}$$

The return of the portfolio, r^p, is simply the sum of realised excess returns multiplied by the corresponding portfolio weights:

$$r_{t+1}^p = \sum_{j=1}^{n} w_{j,t+1} rx_{j,t+1}. \tag{3.9}$$

Notably, the linear portfolio preserves all ratios in size between forecasts. In other words, currencies for which the model predicts extraordinarily large positive or negative risk premia will feature with an extraordinarily large weight in the portfolio. By definition, it is also possible (and likely) that the number of positive portfolio weights is unequal the number of negative portfolio weights.

The rank portfolio, on the other hand, assigns portfolio weights only according to the position of forecasts in a ranking by size, not by the size of the forecast itself. As in Menkhoff et al. (2016) and Asness et al. (2013), we define weights as

$$w_{j,t+1} = c_t\Big(rank(f_{j,t}) - \sum_{j=1}^{N} rank(f_{j,t})/N\Big). \tag{3.10}$$

As with the linear portfolio, the scaling factor c_t is constructed such that the portfolio is long one unit and short one unit of currency. Given the number of currencies in the sample, $c_t = 0.02$, resulting in weights ranging from -0.19 to 0.19 with an interval size of 0.02. By definition, this portfolio will always be long 10 currencies and short 10 currencies.

Finally, we construct a high-minus-low portfolio which only trades the two currencies with the largest and smallest signal. Specifically, the portfolio goes long one unit of the currency with the largest forecast for the risk premium and shorts one unit of the currency with the smallest forecast.

95

Intuitively, the high-minus-low portfolio allows investments to represent very strong views about predicted risk premia to the extent, that "all eggs are being put into one basket". We expect this approach to display very high mean returns, but also a high degree of volatility, as investments are not distributed amongst currencies. The rank portfolio, on the other hand, strongly restricts the model-based formation of a strong view. As mentioned above, the currencies with the strongest signals receive a weight of 0.19 and -0.19, with the strategy ensuring a smooth distribution of investments across all 20 currencies. Consequently, we would expect returns from this approach to exhibit a relatively low volatility. Conceptually, the linear portfolio can be seen as a compromise between the rank and high-minus-low portfolios, where the model is allowed to display conviction whilst still investing in all currencies in the sample.

A property worth pointing out that applies to all three portfolio approaches is that the sign of the portfolio position for a currency does not have to be the same as the sign of the forecast. Even if all currencies are forecast to have positive risk premia in the month to come, the rank portfolio, for example, would short half the currencies, namely those with the smallest appreciation signals. This also applies to the linear and high-minus-low portfolios, although the linear portfolio, as mentioned above, does not necessarily equate the number of long and short positions.

On a general note, portfolios are rebalanced at the end of every month, based on the latest forecast. Trades are executed by buying (selling) 1-month forward contracts and selling (buying) spot at expiry. As such, the portfolios are cost-free.[5]

3.4.4 Benchmark Results

Table 3.2 presents the results for portfolio returns between 2006 and 2016. In particular, it shows annualised mean returns, t-statistics based on Newey and West (1987) standard errors, return volatilities, Sharpe Ratios and hit ratios for linear, rank and high-minus-low portfolios. Hit ratios are defined as the percentage of periods during which the portfolio produces a positive return. Remarkably, all portfolios display large positive mean returns. As expected, the high-minus-low portfolio has the highest annualised mean return at 7.4%. That said, it also displays the highest standard deviation of returns. Whilst the rank portfolio returns have the lowest standard deviation, it is the linear portfolio that has the highest Sharpe Ratio at 0.75. Following the intuitive interpretation above, it seems that the linear portfolio finds the best compromise between allowing the model to show conviction on certain currencies, whilst still optimising the benefits of investing in a portfolio of currencies to smooth returns. Hit ratios for all three portfolios are between 60% and 63%, significantly above the 50% implied by a random walk.

---

[5]For the purpose of our analysis, we disregard margin capital, which under normal circumstances has to be assigned to trades.

	linear	rank	HML
mean	5.05	3.60	7.39
t	[2.76]	[2.36]	[1.96]
ρ	6.59	5.73	13.05
skew	-0.23	-0.27	0.16
kurtosis	0.79	0.72	0.62
Sharpe ratio	0.75	0.62	0.55
hit ratio	63%	60%	61%
biggest monthly gain	6.00	5.31	14.23
biggest monthly loss	-5.62	-5.02	-8.14
maximum drawdown	-11.56	-7.92	-25.78

Table 3.2: This table shows benchmark results for out-of-sample tests of three portfolios from 2006 to 2016. Mean returns, standard deviations and Sharpe ratios are annualised. T-statistics are calculated based on Newey and West (1987) standard errors. The sample includes monthly data from Jan 2005 to Dec 2016. Out-of-sample results are calculated from Jan 2006 to Dec 2016 with an extending window regression.

Note that our portfolio does not profit from the general success of emerging market currencies over the sample period as the portfolio is always short EM currencies to the same extent that it is long other EM currencies.

Figure 3.1 shows cumulative returns as well as drawdowns for the three portfolio strategies. Cumulative returns are calculated as the percent return on one unit invested in the portfolio. Seeing that the portfolios are cost free, no ex-ante investment is required. Consequently, cumulative returns start at 0, as opposed to 100%. Looking at the graph, the most surprising feature is that linear and rank portfolios perform remarkably consistently over the whole sample period—pre-crisis, post-crisis and during the crisis. In fact, there is only one occurence of 5 months and one of 4 months of consecutive negative returns in the linear and rank portfolios (May 2015 to Sep 2015 and Sep 2008 to Dec 2008, respectively). During all other times, consecutive periods of negative returns are at a maximum three months long. In the high-minus-low portfolio, despite incurring larger losses during periods of negative returns, they never exceed 4 months in length. Figure 3.1 again shows very clearly how the commitment to only two currencies in the high-minus-low portfolio comes with a significant increase in downside risk. Drawdowns, calculated as cumulative negative returns since the last maximum in cumulative returns, are far larger for the high-minus-low portfolio compared to the linear and rank portfolios, peaking at over -25%. Linear and rank portfolios, on the other hand, display a very similar pattern with comparably small drawdowns of -12% and -8%, respectively, at worst.

To get a better idea of whether portfolios are able to consistently generate positive returns, we calculate annualised mean returns for every year in the sample period, shown in Figure 3.2. Two observations stand out. Firstly, there are only very few years with negative returns. For the linear portfolio, only three out of eleven years show negative returns. For the rank and

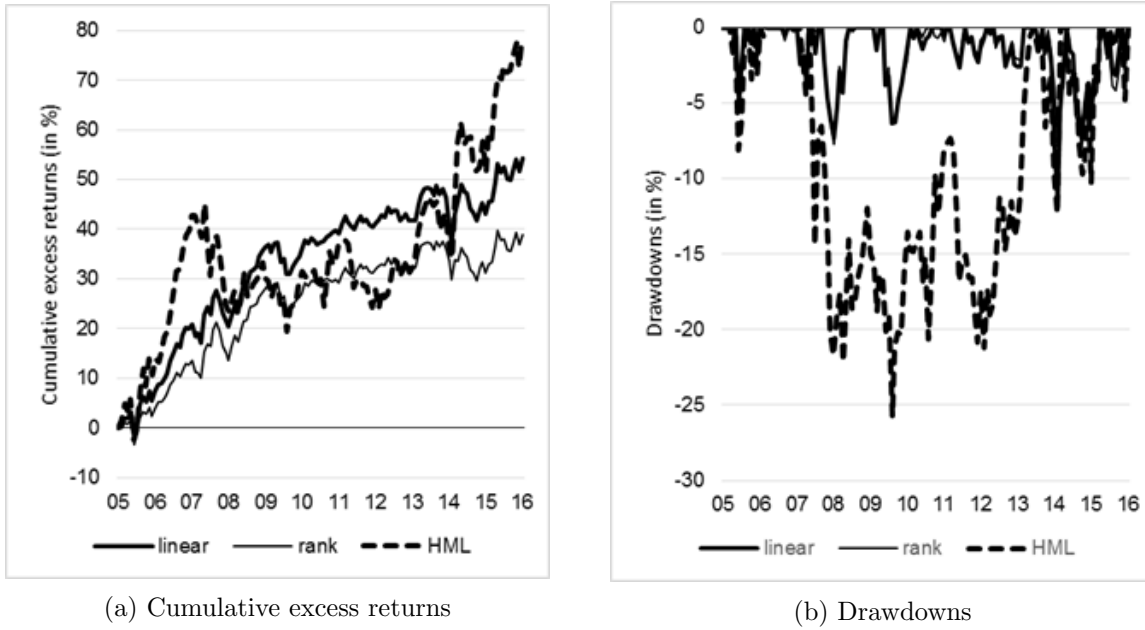| (a) Cumulative excess returns | (b) Drawdowns |

Figure 3.1: Panel (a) shows cumulative portfolio excess returns (in %) for linear, rank and high-minus-low strategies. Panel (b) displays drawdown dynamics of the three portfolios.

high-minus-low portfolios this is even lower at two years each. Secondly, negative returns for both the linear and the rank portfolio are remarkably contained. During 2015, the worst year for the rank portfolio, annualised average returns stayed above -2.5%. 2014, the worst year for the linear portfolio, saw average returns of about -1%. Averaging across all three years of negative returns for the linear portfolio, negative annualised average returns stand at -0.6%. This compares to average returns of 7.4% during years with positive returns.

3.5   Additional Results and Robustness

After a first look at the results, the model seems to be a powerful tool to forecast risk premia in emerging market currencies. But as with all models, there are issues that need to be addressed to establish the robustness of results. We will now cover some of these issues and discuss further avenues for the estimation of risk premia.

Excluding regressors

One of the most appealing factors of the model is its simplicity and straightforward approach. Using only four regressors, however, also raises the question whether the model might be overly dependent on one of them. To get a clearer idea of the influence each regressor has on the forecasts, we replicate all estimations and out-of-sample predictions, excluding one of the regressors at a time.

Table 3.3 presents annualised mean returns, t-statistics based on Newey and West (1987) standard errors, return volatilities, Sharpe Ratios and hit ratios for linear, rank and high-minus-
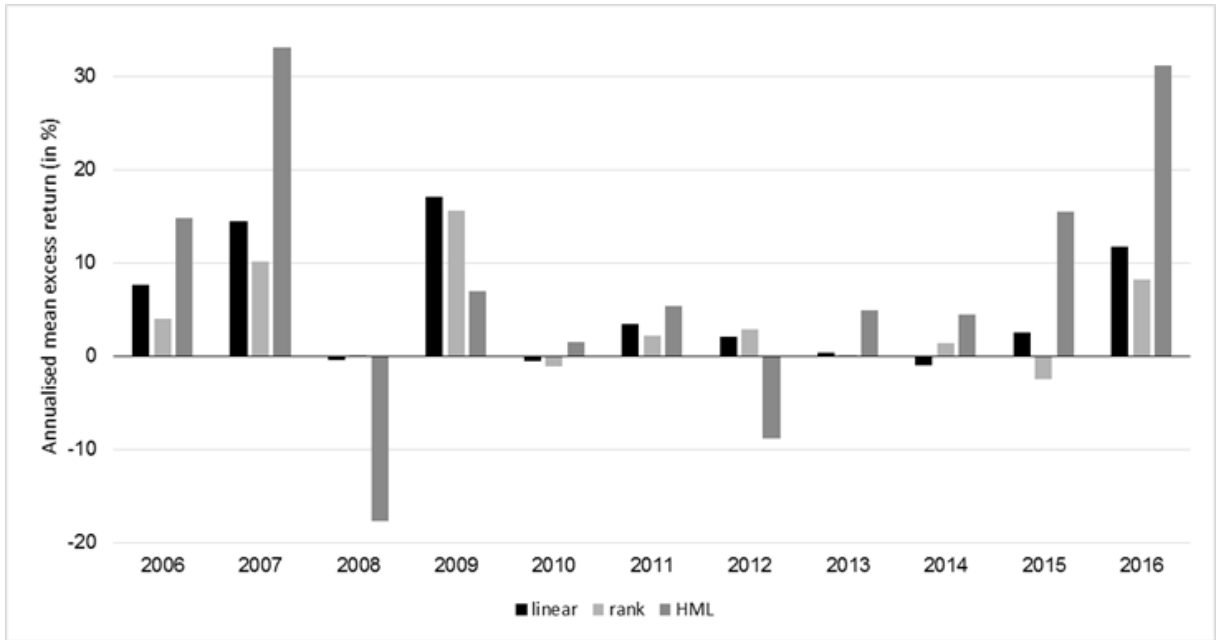
Figure 3.2: Annualised mean excess returns (in %) by year for linear, rank and high-minus-low portfolio strategies.

	excluding *rer*			excluding *ird*		
	linear	rank	HML	linear	rank	HML
mean	4.93	3.63	9.98	3.14	2.56	1.93
t	[2.67]	[2.37]	[2.46]	[2.05]	[1.94]	[0.5]
ρ	6.21	5.33	12.98	5.14	4.68	12.27
SR	0.78	0.67	0.74	0.6	0.54	0.16
hit ratio	61%	60%	62%	55%	57%	53%
	excluding *emom*			excluding *cesi*		
	linear	rank	HML	linear	rank	HML
mean	4.71	4.38	4.89	4.49	2.81	9.01
t	[2.27]	[2.69]	[0.57]	[2.18]	[1.51]	[1.82]
ρ	7.1	5.75	15.45	7.15	6.38	14.6
SR	0.65	0.75	0.31	0.61	0.43	0.59
hit ratio	60%	61%	58%	58%	55%	63%

Table 3.3: This table shows results for out-of-sample tests of linear, rank and high-minus-low portfolios. Compared to the benchmark results in Table 3.2, this table reports results for models where we remove one explanatory variable at a time. Mean returns, standard deviations and Sharpe ratios are annualised. T-statistics are calculated based on Newey and West (1987) standard errors. The sample includes monthly data from Jan 2005 to Dec 2016. Out-of-sample results are calculated from Jan 2006 to Dec 2016.

	linear	rank	HML
mean	3.3	3.72	10.87
t	[1.16]	[1.75]	[1.76]
ρ	8.2	6.6	16.67
skew	-0.71	-0.26	-0.39
kurtosis	1.69	0.91	1.67
Sharpe ratio	0.4	0.55	0.62
hit ratio	59%	59%	63%
biggest monthly gain	6.48	5.51	14.02
biggest monthly loss	-8.25	-6.48	-15.66
maximum drawdown	-21.54	-11.88	-54.38

Table 3.4: This table shows out-of-sample results for carry strategies using linear, rank and high-minus-low portfolios. Mean returns, standard deviations and Sharpe ratios are annualised. T-statistics are calculated based on Newey and West (1987) standard errors. The sample includes monthly data from Jan 2005 to Dec 2016. Out-of-sample results are calculated from Jan 2006 to Dec 2016.

low portfolios. Reassuringly, in no specification do Sharpe ratios for the linear portfolio drop below 0.6, indicating that no single building block of the model drives the results. Interestingly, the Sharpe ratio is even better for the model without real exchange rate strength than for the full model. Whilst this out-of-sample test indicates that the real exchange rate does not have much to contribute to the estimation, we find that the theoretical argument for its inclusion remains overwhelming, even for the short forecast horizon of one month. More generally, from a methodological perspective we would not want to drop regressors on the basis of out-of-sample testing, as this would invalidate the approach. All that said, Sharpe ratios are only one method of evaluating out-of-sample performance, and not an uncontroversial one. Looking for example at hit ratios instead, we see that the full model outperforms the version excluding the real exchange rate.

What About the Random Walk - Is Carry The Way To Go?

Traditionally, the benchmark for success in foreign exchange rate forecasting has been the random walk. Ever since Meese and Rogoff (1983a) suggested that fundamental forecasts do not do better than the random walk, horse races have been run against it. As pointed out in Section 3.3, the random walk hypothesis stipulates that nominal exchange rates are expected to remain unchanged, implying that risk premia correspond one-for-one to interest rate differentials (carry). Consequently, by testing our model against the random walk, we implicitly also test whether it performs better than the classic carry trade. The latter is a relevant question, as many exchange rate models for emerging market currencies seem to perform well simply because they absorb the returns to the carry trade. Table 3.4 presents the results for the model based on carry alone. The standard deviation of returns rises in all three portfolio specifications, whilst mean returns

Sharpe ratios	model	carry	Z-score
linear	0.75	0.4	1.71*
rank	0.62	0.55	0.29
high-minus-low	0.55	0.62	-0.24

Table 3.5: We compare Sharpe ratios of the benchmark model with those of a carry trade using a two-sample version of the Jobson and Korkie (1981) test. The null hypothesis for the three portfolios is that the Sharpe ratios of the benchmark model and of the carry strategy are equal. The last column of the table shows the Z-score difference between the Sharpe ratios, which is approximately normally distributed with mean zero and a standard deviation of one. Only the Z statistic for the linear portfolio is significant (at the 5% level).

fall for both the linear and rank portfolio, leading to a drop in their Sharpe ratios to 0.4 and 0.55, respectively. A rise in mean returns for the high-minus-low portfolio implies a rise in the Sharpe ratio to 0.62. To get a better understanding of the changes in Sharpe ratios, we test whether the Sharpe ratios of the full model are statistically different from the Sharpe ratios of the random walk model. We do so with the two-sample version of the Jobson and Korkie (1981) test. The test is approximately normally distributed with mean zero and a standard deviation of one, and has the null-hypothesis of equal risk-adjusted performance in both models. The results, presented in Table 3.5, show that the linear benchmark model has a significantly higher Sharpe ratio than the carry strategy and thus beats the random walk. Concretely, the null-hypothesis is rejected for the linear model at the 5% level. For the rank and high-minus-low portfolios, however, Sharpe ratios under the benchmark model are not significantly different from carry. This is partly explained by the low power of the test, as noted in Jobson and Korkie (1981) and Jorion (1985).

The test leads us to conclude that carry is not driving the results of our benchmark portfolios.

Excluding Currencies

Similar to the question whether individual regressors drive the results, the question arises whether particular currencies dominate the overall return statistics. To answer this question, we present annualised average return statistics and hit ratios for all currencies over the three portfolio strategies, shown in Table 3.6. Note that, by construction, the annualised mean return of a portfolio is the sum of the mean returns of all 20 currencies displayed in the table. Hit ratios, on the other hand, do not average to the portfolio hit ratio. In fact, whilst the hit ratio for the linear portfolio overall is at 63%, hit ratios for individual currencies only range from between 45% and 62%. It can be inferred that while the portfolio is not particularly good at predicting the direction of each individual currency at every period in time, it is good at assigning higher weights to currencies where it has a high degree of conviction and thus generates high returns when the prediction is correct, compared to relatively small losses when it is wrong. In each of the three portfolios, 8 out of 20 currencies display negative average returns. Figures 3.3, 3.4 and

3.5 graphically show the dispersion in results. Each portfolio strategy suffers from two currencies that significantly breach the 1-standard-deviation band across mean returns. Whilst these are particularly influential in the high-minus-low portfolio, where the Turkish lira (TRY) generates roughly two-thirds of the overall return, the outliers in the linear and rank portfolios are much less consequential. The Turkish lira (TRY) and Polish zloty (PLN), the strongest performers in the linear and rank portfolios, generate less than one-third of the overall portfolio returns.

Note that whilst this analysis sheds light on the consistency of returns across different parts of the portfolio, we cannot infer what returns would be were we to exclude one of the currencies. The exclusion of a currency changes weights for all other currencies in the portfolio, leading to different return patterns across the sample.

Table 3.6: The table shows out-of-sample results of the benchmark portfolios, by currency. Mean returns are annualised. Note that by construction, the annualised mean return of a portfolio is the sum of all 20 currency mean returns displayed in the table. Hit ratios, on the other hand, do not average to the portfolio hit ratio. The sample includes monthly data from 2006M1 to 2016M12.

|	BRL	CLP	COP	CZK	HUF	INR	IDR	ILS	MYR	MXN	PEN	PHP	PLN	RUB	ZAR	KRW	THB	TRY	SGD	TWD
linear portfolio																				
mean	-0.08	-0.02	0.54	-0.29	0.75	-0.13	0.49	-0.11	-0.03	-0.02	0.15	0.52	0.81	0.20	-0.68	0.45	0.35	1.65	0.04	0.37
hit ratio	59%	50%	54%	45%	61%	48%	53%	48%	49%	45%	48%	48%	55%	46%	49%	58%	54%	59%	50%	56%
rank portfolio																				
mean	-0.28	-0.18	0.43	-0.16	0.63	-0.04	0.37	-0.09	-0.08	-0.04	0.13	0.60	0.89	0.07	-0.48	0.32	0.33	0.68	0.09	0.37
hit ratio	58%	51%	53%	44%	62%	50%	53%	50%	50%	47%	46%	51%	59%	48%	50%	57%	54%	58%	50%	60%
high-minus-low portfolio																				
mean	-1.04	1.05	1.09	-3.04	0.70	-0.53	0.00	-0.14	-0.18	0.27	-0.02	1.15	1.03	0.56	-0.82	0.38	-0.22	5.07	0.61	1.45
hit ratio	31%	100%	62%	29%	67%	25%	50%	38%	50%	50%	57%	69%	100%	58%	40%	50%	43%	70%	48%	80%

103
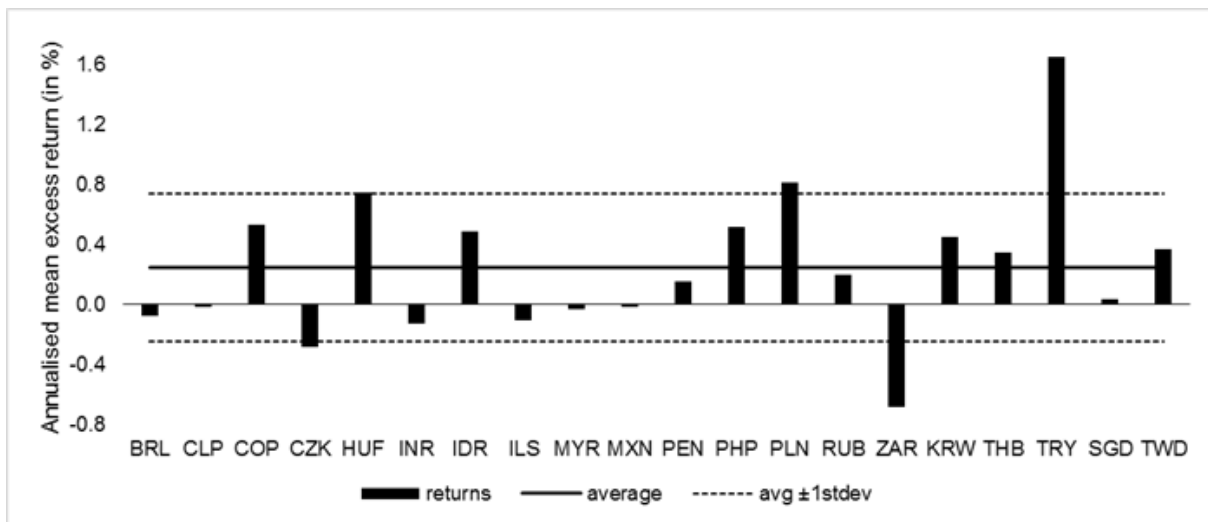
Figure 3.3: Annualised mean excess returns (in %) by currency for linear portfolio strategy.

Decomposing The Signal

In constructing the model, we chose to allow for changing coefficients over time. This gives the model the opportunity to adapt to changing circumstances and relationships. It raises the question, however, whether the model might behave in an erratic way and react excessively to short-term changes in the market environment. To gain a better understanding on how exactly parameters behave, we dive deeper into two aspects of the model. Firstly, we decompose the signal to measure the influence of regressors and how this changed over time. Secondly, we look at how the estimated coefficients have changed over the sample period.

Figure 3.6 portrays the influence each regressor has on the overall signal in each period. To construct the series, we multiply each coefficient with the standard deviation of the corresponding variable and compare the absolute value to that of other coefficients as a percentage share. A number of observations are noteworthy. Firstly, the model saw a lot of change between 2006 and 2009, but has been remarkably stable since then. Growing steadily over the first years of the sample, equity momentum has developed to become the dominating element of forecasts. At the same time, carry lost some of its importance after 2009 and is roughly of equal importance as the real exchange rate and economic surprises.

Figure 3.7 looks in more detail at the behaviour of individual coefficients. Each line in the graph represents the product of a variable's regression coefficient and the variable's standard deviation. Overall, the graph confirms the picture of broadly stable relationships, especially after 2009. That said, a paradigm shift seems to take place when it comes to the influence of equity momentum (long dashed line). Starting off with a significant negative coefficient, i.e. representing a reversal element where strong equity performance is followed by negative risk premia, the coefficient turns positive very suddenly in mid-2008, introducing the notion of momentum in the place of reversal. All other coefficients have the sign that theory would
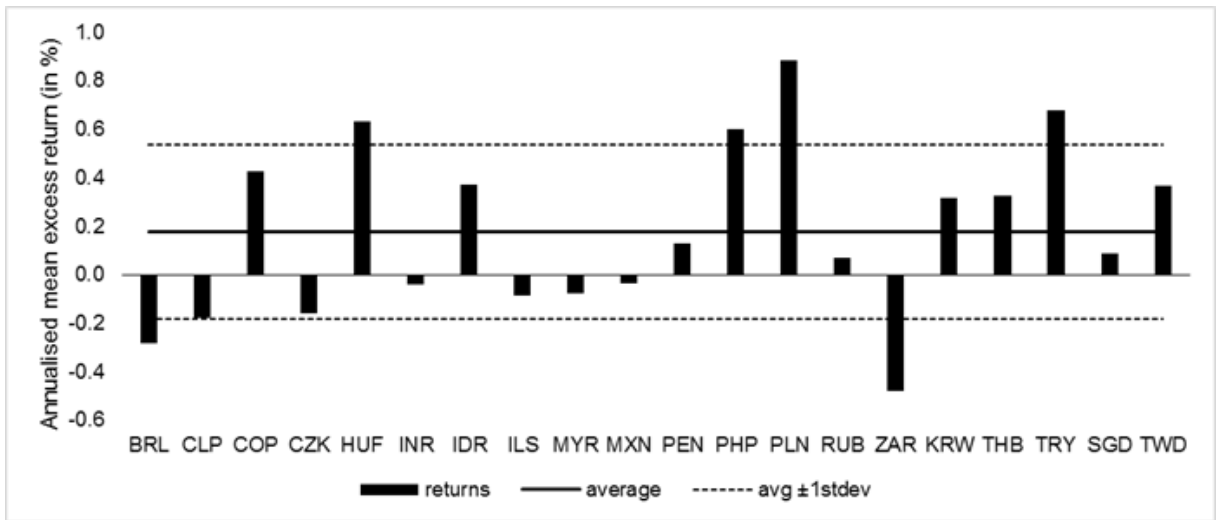
Figure 3.4: Annualised mean excess returns (in %) by currency for rank portfolio strategy.
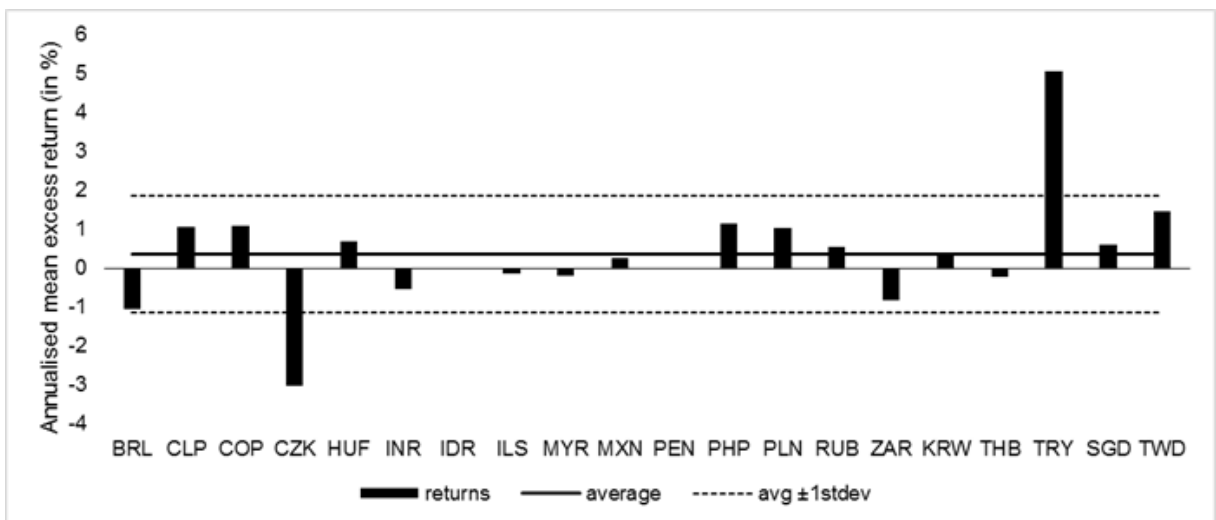

Figure 3.5: Annualised mean excess returns (in %) by currency for high-minus-low portfolio strategy.
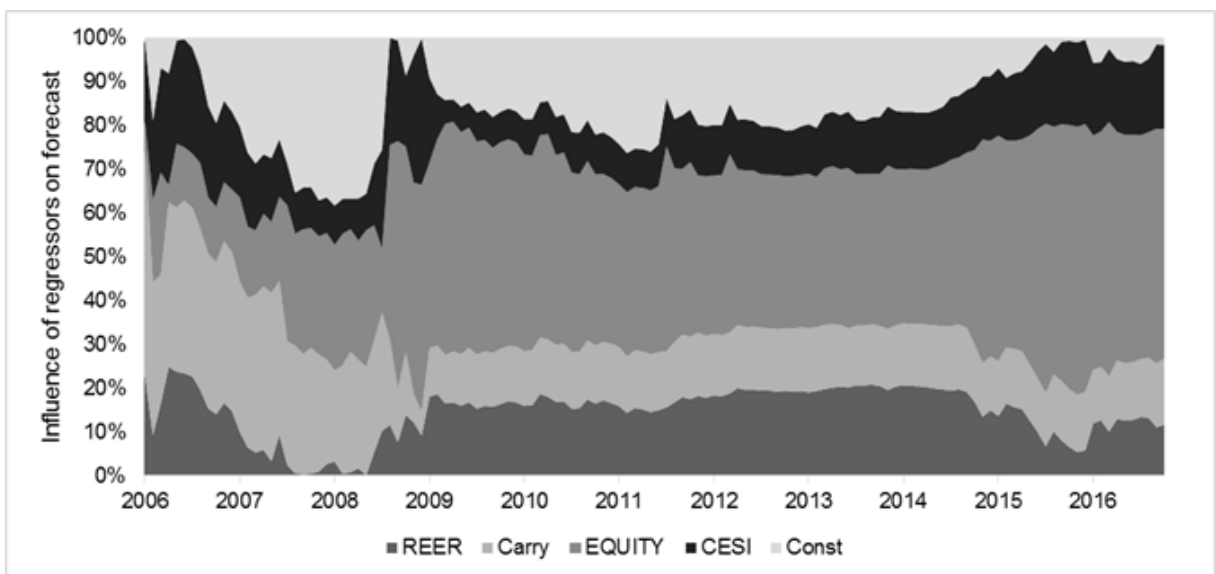

Figure 3.6: Influence of explanatory variables on overall forecast signal, constructed as the absolute value of product of regression coefficient and standard deviation of the corresponding variable.
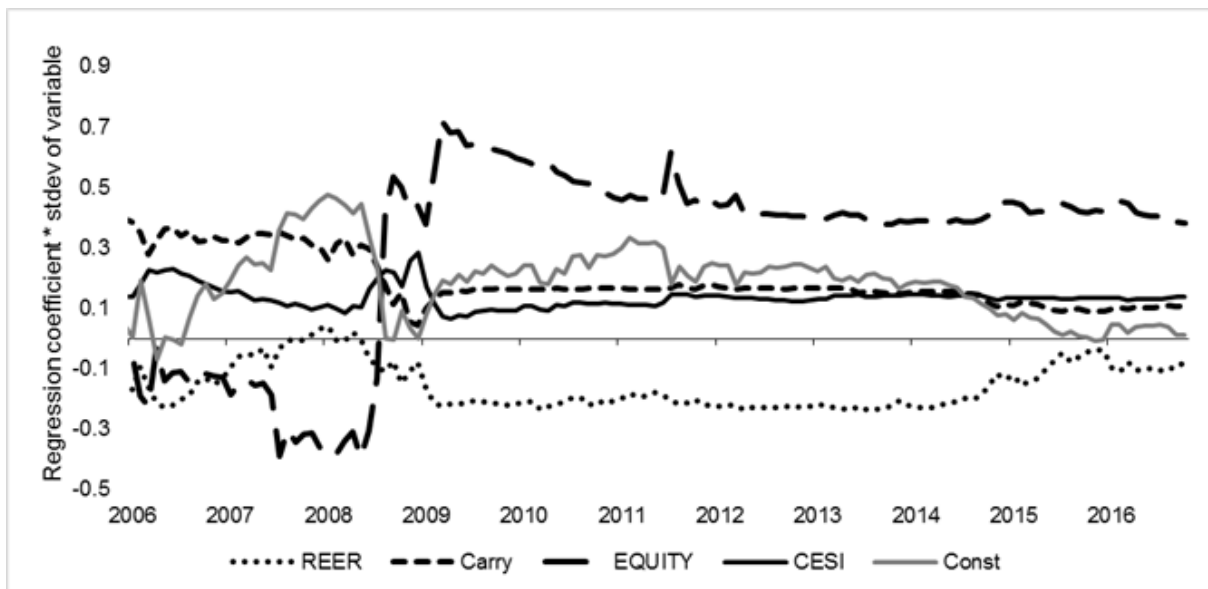
Figure 3.7: Influence of explanatory variables on overall forecast signal, constructed as the product of a variable's regression coefficient and standard deviation of the variable.

suggest: real exchange rate strength has a negative effect on risk premia, while it is positive for carry and economic surprises.

Implied USD Exposure

By construction, all portfolio strategies presented above have a neutral position to the US dollar. While the dollar serves as the hypothetical home currency against which all returns are measured, we never enter either a long or a short position in USD. That said, it is well known that some emerging market currencies have a higher correlation to the USD than others. This implies that, in practice, our portfolios could have an indirect exposure to USD movements, whether intended or not. In other words, the portfolio can profit from exposing itself to the USD by opening positions in EM currencies that are themselves highly correlated with the dollar.

Examining this question further, we first calculate 12-month rolling correlations between 1-month nominal exchange rate changes and the 1-month return in the broad trade weighted USD index. In a second step, we weight these correlations with the portfolio weights for each currency and sum them up to create a measure of portfolio exposure to the dollar. Figure 3.8 shows this measure of dollar exposure for the linear, rank and high-minus-low portoflios. The chart also shows the history of the USD index. As can be seen, the portfolios are indeed exposed to the USD, at times profiting and at other times losing out from movements in the dollar. An interesting avenue for future investigation would be to examine whether these exposures can be systematically exploited to increase portfolio returns.
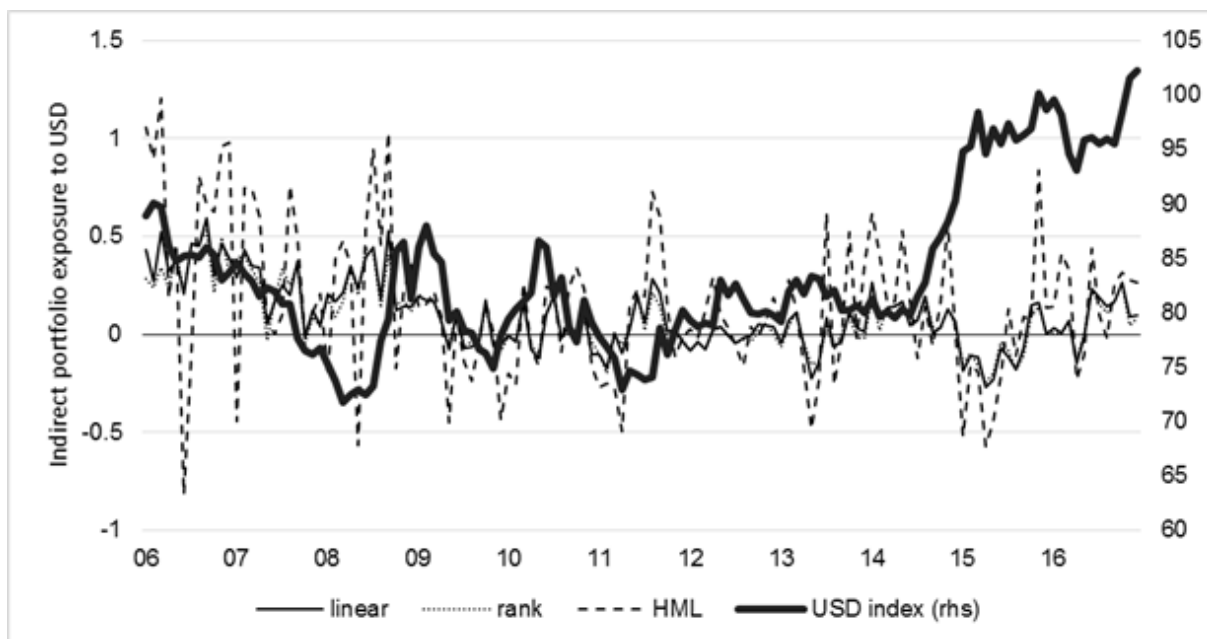
Figure 3.8: Indirect portfolio exposure to USD index, calculated as the weighted sum of rolling correlations between EM currency returns and USD index returns. We use Bloomberg's DXY Index for the USD, plotted on the right side of the chart.

Transaction Costs

Finally, an important question to ask is in how far transaction costs reduce realised out-of-sample portfolio returns. We use data on bid-ask spreads available on Bloomberg and from BofA Merrill Lynch (2014) to recalculate the statistics from Table 3.2, adjusting returns for transaction costs. Average transaction costs across all currencies in the sample are 3.8 basis points,[6] implying that the monthly return to an investment of 1 unit of currency is reduced by approximately 0.038 percentage points. Note that these calculations can only serve as an approximation, as transaction costs change over time. Moreover, the size of bid-ask spreads in practice depends on the amount traded.

Table 3.7 presents the recalculated results, alongside the original numbers. As can be seen, Sharpe Ratios fall by between 0.04 and 0.08, but remain large and significant.

More generally, every reduction in Sharpe Ratios naturally reduces the willingness of market participants to engage in a certain strategy. As is the case with most proposed currency strategies, significant returns might simply exist because the costs of arbitrage are too high. Whether taking transaction costs into account or not, this possibility cannot be excluded ex ante.

---

[6]Each position requires payment of the complete bid-ask spread, as half has to be paid when entering a position and the other half again when exiting a position.

	no transaction cost			incl. transaction cost		
portfolio	linear	rank	HML	linear	rank	HML
mean	5.05	3.60	7.39	4.57	3.13	6.91
t	[2.76]	[2.36]	[1.96]	[2.51]	[2.05]	[1.84]
ρ	6.59	5.73	13.05	6.59	5.73	13.05
Sharpe ratio	0.75	0.62	0.55	0.68	0.54	0.51

Table 3.7: This table shows results for out-of-sample tests of the three benchmark portfolios, adjusted for transaction costs. Mean returns, standard deviations and Sharpe ratios are annualised. T-statistics are calculated based on Newey and West (1987) standard errors. The sample includes monthly data from 2006M1 to 2016M12.

3.6 Conclusion

Bridging various strands of the existing literature on exchange rate forecasting, this paper has proposed a novel approach to estimating currency risk premia. The analysis has produced three main results.

Firstly, the paper has shown that combining a measure of currency value in the form of real exchange rate strength with measures for market sentiment allows for the creation of powerful forecasts over a short horizon. Whilst most currency forecasts focus on horizons of multiple years and, at an extreme, manage to forecast currencies over one quarter, the model proposed in this paper predicts excess returns over only one month. In doing so, we improve upon the widely established result that excess returns cannot be consistently forecast for such a short time horizon.

Secondly, the model outperforms a carry trading strategy both in terms of the size of returns, as well as their consistency. By design, the portfolio approach from our benchmark model can be replicated for a carry trading strategy, allowing results between carry and the benchmark model to be directly compared. As the carry strategy in our model is by definition identical to a random walk model, we thereby also beat the latter—widely considered to be the quintessential benchmark in short term currency forecasting.

Thirdly, we show that results hold up to the inclusion of transaction costs. While returns are smaller across the portfolio, they are still significant and economically large, allowing investors to systematically generate positive excess returns in emerging market currencies.

All results have been established out of sample and further checks establish their robustness across time periods, currencies included and explanatory variables.

Methodologically, the model takes a clear stance when it comes to arbitrary and heterogeneous models of currency forecasting. For one, the model only includes a limited number of explanatory variables that are all identified by the existing literature as highly significant for forecasting risk premia. Moreover, the model applies the same data and an identical structure across all currencies in the sample, providing a generalised method that is independent of the

chosen currency and time period. Using an OLS estimation technique furthermore allows a straightforward decomposition and interpretation of forecasting signals, providing information beyond the numerical output itself.

Forecasting currency returns remains an extremely difficult challenge, where results are often invalidated as the market learns to incorporate them. In this environment, progress in the literature is often met with further increasing gaps, leaving a lot of room for improvement in our understanding of underlying mechanisms. As currency forecasting will never be a hard science, it remains tantamount that models not only produce a mechanical output, but also provide the market observer with information as to the origin and composition of signals. In other words, black boxes are to be avoided. This paper has taken a significant step in producing a model that is general, yet simple enough to provide information far beyond each currency's numerical forecast and, in doing so, filled gaps in the literature between long-term and short-term forecasting, as well as between fundamental and behavioural forecasting. That said, a lot of work remains to be done to better understand and predict exchange rate movements. Whilst further theoretical contributions are needed that build models incorporating the empirical evidence, the wider availability of data will make it possible to also further extend empirical analyses of currencies. One direct extension to the paper will be an analysis of the proposed model for developed market currencies. Another promising and highly relevant direction of research will be to create better measures of market sentiment and uncertainty on a country by country level. These remain important avenues for future research.

Bibliography

Asness, C. S., Moskowitz, T. J., and Pedersen, L. H. (2013). Value and momentum everywhere. *The Journal of Finance*, 68(3):929–985.

Bank for International Settlements (2016). Foreign exchange turnover in April 2016. Triennial Central Bank Survey, Monetary and Economic Department.

Barroso, P. and Santa-Clara, P. (2015). Beyond the carry trade: Optimal currency portfolios. *Journal of Financial and Quantitative Analysis*, 50(5):1037–1056.

BofA Merrill Lynch (2014). 2014 EM Local Markets Guide. GEMs Paper 18, GEM Fixed Income Strategy & Economics.

Burnside, C., Eichenbaum, M., Kleshchelski, I., and Rebelo, S. (2011a). Do peso problems explain the returns to the carry trade? *Review of Financial Studies*, 24(3):853–891.

Burnside, C., Eichenbaum, M., and Rebelo, S. (2011b). Carry trade and momentum in currency markets. *Annual Review of Financial Economics*, (3):511–35.

Cenedese, G., Payne, R., Sarno, L., and Valente, G. (2015). What do stock markets tell us about exchange rates? *Review of Finance*, 20(3):1045.

Clark, P. B. and MacDonald, R. (1999). *Exchange Rates and Economic Fundamentals: A Methodological Comparison of Beers and Feers*. Springer Netherlands, Dordrecht.

Driver, R. L. and Westaway, P. F. (2004). Concepts of equilibrium exchange rates. Working Paper no. 248, Bank of England.

Engel, C. (2016). Exchange rates, interest rates, and the risk premium. *American Economic Review*, 106(2):436–74.

Engel, C., Mark, N. C., and West, K. D. (2008). *Exchange Rate Models Are Not As Bad As You Think*, pages 381–441. University of Chicago Press.

Engel, C., Mark, N. C., and West, K. D. (2012). Factor model forecasts of exchange rates. Working Paper 18382, National Bureau of Economic Research.

Engel, C. and West, K. D. (2010). Global interest rates, currency returns, and the real value of the dollar. *American Economic Review*, 100(2):562–67.

Flood, R., Marion, N., and Albornoz, J. Y. (2012). A common framework for thinking about currency crises. In *Handbook of Exchange Rates*. Wiley & Sons, Inc.

Froot, K. A. and Ramadorai, T. (2005). Currency returns, intrinsic value, and institutional-investor flows. *The Journal of Finance*, 60(3):1535–1566.

James, J., Marsh, I. W., and Sarno, L., editors (2012). *Handbook of Exchange Rates*. Wiley Handbooks in Financial Engineering and Econometrics. John Wiley & Sons, Inc.

Jobson, J. D. and Korkie, B. M. (1981). Performance hypothesis testing with the sharpe and treynor measures. *The Journal of Finance*, 36(4):889–908.

Jorda, O. and Taylor, A. M. (2009). The carry trade and fundamentals: Nothing to fear but feer itself. Working Paper 15518, National Bureau of Economic Research.

Jorion, P. (1985). International portfolio diversification with estimation risk. *The Journal of Business*, 58(3):259–278.

Koijen, R. S. J., Moskowitz, T. J., Pedersen, L. H., and Vrugt, E. B. (2016). Carry. Fama-Miller Working Paper.

Lustig, H., Roussanov, N., and Verdelhan, A. (2011). Common risk factors in currency markets. *Review of Financial Studies*, 24(11):3731–3777.

Lustig, H. and Verdelhan, A. (2007). The cross section of foreign currency risk premia and consumption growth risk. *American Economic Review*, 97(1):89–117.

Mark, N. (1995). Exchange rates and fundamentals: Evidence on long-horizon predictability. *American Economic Review*, 85(1):201–18.

Mark, N. C. and Sul, D. (2012). When are pooled panel-data regression forecasts of exchange rates more accurate than the time-series regression forecasts? In *Handbook of Exchange Rates*. Wiley & Sons, Inc.

Meese, R. and Rogoff, K. (1983a). Empirical Exchange Rate Models of the 1970s: Do They Fit Out of Sample? *Journal of International Economics*.

Meese, R. and Rogoff, K. (1983b). *Exchange Rates and International Macroeconomics*, chapter The Out-of-Sample Failure of Empirical Exchange Rate Models: Sampling Error or Misspecification?, pages 67–112. University of Chicago Press.

Meese, R. and Singleton, K. J. (1982). On unit roots and the empirical modeling of exchange rates. *The Journal of Finance*, 37(4):1029–1035.

Menkhoff, L., Sarno, L., Schmeling, M., and Schrimpf, A. (2011). Currency momentum strategies. Working paper.

Menkhoff, L., Sarno, L., Schmeling, M., and Schrimpf, A. (2016). Currency value. *Review of Financial Studies*.

Menkhoff, L. and Taylor, M. P. (2007). The obstinate passion of foreign exchange professionals: Technical analysis. *Journal of Economic Literature*, 45(4):936–972.

Neely, C. J. and Weller, P. A. (2012). Technical analysis in the foreign exchange market. In *Handbook of Exchange Rates*. Wiley & Sons, Inc.

Newey, W. and West, K. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–08.

Obstfeld, M. and Rogoff, K. (1996). *Foundations of International Macroeconomics*. The MIT Press, 1st edition.

Pring, M. J. (1991). *Technical analysis explained*. McGraw-Hill.

Rabanal, P. and Rubio-Ramirez, J. F. (2015). Can international macroeconomic models explain low-frequency movements of real exchange rates? *Journal of International Economics*, 96(1):199 – 211.

Rogoff, K. (2001). The failure of empirical exchange rate models: No longer new but still true. *Economic Policy Web Essay*.

Rogoff, K. (2008). Comment on "exchange rate models are not as bad as you think". In Acemoglu, D., Rogoff, K., and Woodford, M., editors, *NBER Macroeconomics Annual 2007*, volume 22. University of Chicago Press.

Sarno, L. and Taylor, M. P. (2002). *The Economics of Exchange Rates*. Cambridge University Press.

The Economist (2013). Carry on trading: why nominal interest-rate differentials are important to currency markets. August 10th 2013.

Trivedi, K. and Ozerov, M. (2016). Finding fair value in EM FX. Global Economics Paper 227, Goldman Sachs International.