



BANK OF ENGLAND

Staff Working Paper No. 646

What drives business investment in the United Kingdom? Results from a firm-level VAR approach

Marko Melolinna

February 2017

Staff Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Bank of England or to state Bank of England policy. This paper should therefore not be reported as representing the views of the Bank of England or members of the Monetary Policy Committee, Financial Policy Committee or Prudential Regulation Authority Board.



BANK OF ENGLAND

Staff Working Paper No. 646

What drives business investment in the United Kingdom? Results from a firm-level VAR approach

Marko Melolinna⁽¹⁾

Abstract

This paper studies the effects of macroeconomic shocks on business investment in the United Kingdom by filtering a large UK firm-level based dataset of financial accounts into macro-level proxy indicators, and then using these indicators in a Bayesian vector autoregression framework to analyse these effects. The analysis combines micro-level data with macro-level analysis in a unique way, and brings up several interesting empirical results. Supply shocks have tended to have been more persistent and more important than demand shocks in explaining UK investment dynamics over the past fifteen years, and their importance appears to have increased since the financial crisis. Furthermore, shocks to the cost of capital, and uncertainties related to it, have generally been more important for firms in sectors with higher indebtedness, whereas corporate governance issues as measured by dividend payments and share buybacks do not appear to have been a major driver of investment.

Key words: Business cycle, micro data, vector autoregression, sign restrictions, time-varying parameters.

JEL classification: C11, C32, D21, E32, E52.

(1) Bank of England. Email: marko.melolinna@bankofengland.co.uk

The views expressed here are solely my own, and do not necessarily reflect those of the Bank of England. I thank Konstantinos Theodoridis for providing his code and George Kapetanios, Patrick Schneider, Garry Young and an anonymous referee for their helpful comments.

Information on the Bank's working paper series can be found at www.bankofengland.co.uk/research/Pages/workingpapers/default.aspx

Publications and Design Team, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 email publications@bankofengland.co.uk

1. Introduction

Depending on the modelling framework used, business investment would typically be expected to be affected by a number of factors, including the cost of capital, value of capital stocks, price of investment goods and general macroeconomic uncertainty. However, in practice, modelling business investment dynamics is challenging. The data on investment are often volatile, and traditionally established relationships between investment and its drivers (like lower cost of capital leading to higher investment) do not always appear to hold, especially in the post-financial crisis world (see, for example, Gilchrist and Zakrajsek (2007) and Gertler and Karadi (2011)). Given that ultimately investment is driven by firm-level decisions, research has recently explored the potential advantages of using firm-level data for analysing investment dynamics. This also allows for heterogeneity between firms to be taken into account as a factor adding variation to the analysis. For example, work on the UK has studied the relationship between company cash-flows, profitability and investment with firm-level data (see Bond et al. (2004) and Farrant et al. (2013)).

I take a slightly different approach to using firm-level data to the past literature. One problem typically encountered with micro data is that it is relatively noisy, especially when dealing with quarterly financial account data. Transforming micro level data into a form that is applicable to macro level analysis is typically also challenging. The aim of this paper is to overcome these problems related to noisy micro data by studying the transmission of macroeconomic shocks to business investment in the UK, using aggregate-level financial account data filtered from firm-level financial account variables. Of particular interest is the transmission of macroeconomic shocks to firm-level investment decisions. I also use time-varying parameter time series methods to study investment dynamics around the financial crisis period.

The implicit suggestion I am making is that rather than using aggregate national accounts data, potentially useful information about macro-level links between key variables can be inferred by using a large number of time series reflecting firm-level behaviour. There are a number of reasons for this to be the case. First, using firm-level data is more in line with micro-founded models based on a firm's profit-maximisation problem (see, for example, Gilchrist and Zakrajsek (2014)), which can reveal different dynamics compared to aggregate level data. Second, basing the modelling framework on financial account data potentially allows for studying much richer inter-dependencies between specific financial account items and firm capital expenditure decisions (including linkages between financing and investment decisions) than aggregate data, as well as cutting the data based on firm characteristics. In the current study, I provide some examples of this type of analysis based on firm sector, size and corporate governance related issues. Third, firm-level data allows for a more genuine definition of supply shocks, as firms' cost dynamics can be measured directly from their financial accounts, rather than relying on aggregate level price indices (where typically versions of consumer

prices are used). Finally, financial account data is much less prone to revisions than national accounts data, potentially allowing for more timely analysis of recent investment dynamics.

Methodologically, the main tools used are factor models and sign-restricted Bayesian vector autoregression (VAR) models. The factor model framework allows for missing observations to elicit proxy indicator time series for certain key series in different sub-sectors and for different firm sizes. The idea is to use financial account items as proxies for headline macroeconomic variables to pick the key series; investment can be proxied by capital expenditure, inflation by cost of goods sold and GDP by operating income.¹ The indicator time series, formed on the basis of the key proxy series, are plugged into a Bayesian macro/monetary policy VAR. Impulse responses and other analytics are then computed to analyse the effects of conventional macroeconomic demand, supply and interest rate shocks (see, for example, Christiano et al. (1996), Peersman and Smets (2001) and Uhlig (2005)) on business investment.

The VAR model framework is also similar in spirit to the factor-augmented VAR (FAVAR) introduced by Bernanke et al. (2005). FAVAR models typically include the monetary policy rate as an “observable” variable, while “unobservable” variables (in the current case, GDP, investment and inflation), are represented by indicator series, or factors, derived from a large number of series. The rationale behind FAVAR models is that because all macroeconomic data, including national accounts data, are noisy, a more appropriate signal for dynamics in key macroeconomic phenomena can be gleaned from a large dataset representing different aspects of those phenomena. While the underlying series in Bernanke et al. (2005) refer to macro-level series and those in the current study to firm-level series, the logic behind the modelling framework used is the same.

The main contribution of the paper is to present a unique way of combining micro- and macro-level data for analysing investment with various time series methods. There are a number of empirical results worth highlighting. First, supply shocks have are estimated to have been more persistent and more important than demand shocks in explaining investment dynamics in the UK over the past 15 years, and their importance appears to have increased since the financial crisis. Second, there are sectoral differences in investment dynamics, especially in terms of explaining variation in investment. Third, shocks to the cost of capital, and uncertainties related to it, are estimated to have generally been more important in explaining investment dynamics for firms in sectors with higher indebtedness. Fourth, demand shocks are estimated to have been more important for large firms than SMEs, whereas the more heterogeneous nature of SMEs probably accounts for the unexplained residual component being an important investment driver for these firms. A modified version of the model

¹ Firm operating income is not the same concept as GDP, which measures the gross value added. However, for practical purposes, and when no other options exist, it is common practice in the literature to use operating income as a proxy for GDP (see, for example, Bond et al. (2003)). The empirical result on the correlation between the two series reported below also supports this choice.

also suggests that recent increases in dividend payments and share buybacks do not appear to have been a major driver in “crowding out” investment.

The paper is organised as follows. Section 2 introduces the modelling framework. Section 3 describes the data used in the analysis. Section 4 presents the results and section 5 concludes. Most of the technical details are relegated to appendices.

2. The econometric specification

The VAR modelling approach used in the study involves two stages. First, I need to filter the firm-level micro data into aggregate level indicators of macroeconomic phenomena. Second, I use these indicators in a traditional sign-restricted VAR framework to examine the effects of macroeconomic shocks on firms’ investment dynamics. These two stages are described in the following subsections.

2.1 Filtering the firm-level data

The firm-level micro data needs to be transformed into macro-level indicators for the relevant variables used in the model. Normally, it is common to use techniques like principal component (PC) analysis to achieve this. However, for the current study, this is not an option; the firm-level dataset is relatively large and has a significant number of missing observations – a circumstance unsuitable for PC analysis (see section 3 for more specific descriptions of the data).

There are methods in the literature to overcome this problem. In particular, a more robust filtering method is presented by Giannone et al. (2008), which allows for a common factor to be extracted for large datasets with missing observations.² The basic idea of the methodology is to use Kalman filtering methods to provide consistent estimates of common factors, even when some observations within the dataset are missing.

In this way, it is possible to estimate underlying common factors for the key variables of the model (capex, operating income and total costs). The following choices and assumptions made when using the filter models are worth highlighting:

- i) The data needs to be in stationary format. Given the nature of the dataset and the volatility of individual data series, year-on-year percentage changes of the series are used in the models, as using either levels (often non-stationary and not cointegrated) or quarterly changes (often too noisy) are not viable options.
- ii) It is assumed that one common factor is enough to capture the underlying dynamics of the variable. While this may seem controversial, the model is relatively robust to the inclusion of an additional factor (benchmark case in the empirical analysis of Giannone et

² See Appendix 2 for details of the methodology.

al. (2008)), and in this 2-factor case, the resulting dynamics of the first common factor are very closely correlated with the corresponding macro series (see next section).

- iii) Only series for which at least two-thirds of the observations over the sample period are available are included in the factor models. This is also the benchmark assumption in Giannone et al. (2008), and again, results are robust to changes in this ratio.

2.2 VAR framework

Consider a standard vector autoregression model of order p (VAR(p)) in reduced form (see, for example, Lutkepohl (2005)):

$$y_t = v + A_1 y_{t-1} + \dots + A_p y_{t-p} + u_t \quad (1)$$

where y_t is a ($K \times 1$) vector of dependent variables, v is a ($K \times 1$) constant term vector, u_t is an i.i.d. error term and $A_1 \dots A_p$ are ($K \times K$) coefficient matrices. This model is used both at the whole economy as well as sectoral level (as described below).

The benchmark model includes four variables ($K=4$), based on the data described in the next section. The model is kept purposefully parsimonious, and the choice of variables has its theoretical foundations on a neo-classical interpretation of factors affecting investment (for the original framework, see Jorgenson (1963) and for a concise summary of the literature, see Baumann and Price (2007)). An indicator measure of real capital expenditure (a proxy for business investment), operating income (a proxy for GDP), total costs (a proxy for cost pressures) and a cost of capital measure. An example of adding a variable (i.e., $K=5$) is presented in Section 4.2. The sample is from 2000Q1 to 2014Q4, which is relatively short, but is dictated by the availability of the firm-level data. As is conventional in quarterly VAR models, $p=4$.

The analysis is carried out with a Bayesian version of the VAR using sign restrictions.³ Given the nature of the firm-level data used in the analysis, the Bayesian treatment of the parameters as random variables with potentially different distributions seems appropriate. A Minnesota prior (see Litterman (1986) for details) with 2,000 iterations (burn-in of 1,000) is used, although the results are relatively robust to using other priors. The hyperparameters of the model are chosen with a grid search, which optimises the marginal likelihood from all combinations of the parameters.

To identify of shocks in the model, I use a sign restriction strategy⁴ introduced by Uhlig (2005) and refined by Rubio-Ramirez et al. (2010)). This type of identification strategy is popular in the

³ The estimation is done with the Bayesian Estimation, Analysis and Regression (BEAR) toolbox for Matlab, developed at the European Central Bank.

⁴ For more details of the sign restriction strategy see Appendix 4.

literature⁵, and is often preferred to a more arbitrary Choleski type ordering. The sign restriction methodology also has advantages in the current study. First, it allows for identification of macroeconomic shocks (rather than just shocks to individual variables in a Choleski setup). Second, Choleski ordering may not be appropriate; for example, it is not *a priori* clear why firms would not respond to changes in cost of capital by changing their investment plans within the same quarter. Nevertheless, as a robustness check, some of the results are also presented based on the Choleski ordering.

The sign restrictions imposed to identify the shocks in the model are presented in Table 1. There are four types of shocks in the model; demand, supply, cost of capital and “residual”. In a traditional demand/supply framework, demand shocks can be thought of as shocks that move the demand curve to the right, causing an increase in both activity – including investment – and prices. Supply shocks (typically cost shocks, like oil supply disruptions), on the other hand, move the supply curve to the left, leading to an increase in prices but a decrease in activity. As the reaction of monetary policy and the cost of capital to this type of a shock is not *ex ante* obvious, this restriction is not imposed for the supply shock. In addition to the demand and supply shocks, cost of capital shocks are also included, with the conventional restrictions that in the short term, lower interest rates will lead to an increase in prices and in real activity (or at least not to a decrease in prices and activity). The last, “residual” shock can be assumed to include mainly expectational shocks as well as factors affecting investment not captured elsewhere in this relatively simple linear framework.

Table 1: Sign restrictions with pure sign restriction approach

Shock\variable	Real capital expenditure	Costs	Real operating income	Cost of capital
Demand shock	+	+	+	+
Supply shock	-	+	-	
Cost of capital shock	-	-	-	+
“Residual” shock	+			

Note: empty cell indicates the sign is not restricted. All restrictions also include a zero response.

The model also includes three annual dummy variables for years 2007 to 2009, covering the years around the height of the financial crisis. Ideally, to fully capture the effects of the shocks during the crisis period, we would want to exclude (at least some of) the dummies. However, this turns out not to

⁵ The literature is too extensive to be referenced in any detail here, but for examples in different macroeconomic setups, see Canova and Paustian (2011), Duchi and Elbourne (2016), Melolinna (2012) and Peersman (2005).

be possible, as some of the models become unstable. This problem is caused by the fluctuations in year-on-year growth around the financial crisis, which cannot be avoided due to way the micro data is constructed for the models. This is one of the trade-offs of using this type of data. Nevertheless, given that the stability of the models is preserved, this should not pose a problem for interpreting the results.

3. Data

The basic idea in the current study is to combine firm-level micro data with whole economy level macro data in a unique way to combat the known problems of micro-level data. Irrespective of its source, micro data is noisy and typically requires considerable cleaning and manipulation before it can be used for macro level analysis. This cleaning process often relies on the discretion of the researcher and is not robust to different choices. The results can, for example, be affected by the choice of imputation methods for missing observations in an unbalanced panel.

As described below, I aim to side-step these issues related to traditional micro-data analysis by filtering a large number of micro series into a single proxy indicator and then comparing the resulting series with corresponding macro level data real GDP, business investment and costs. The way this proceeds is taking the relevant micro data for each variable (e.g., capital expenditure), filtering all the firm-level time series with the methods described above in Section 2.1, and using the resulting factor (F_t in equations (A7)-(A8) in Appendix 2) in the VAR framework described in Section 2.2.

Combining micro and macro level data in this way is unconventional in the literature, but I believe it brings clear benefits for the current analysis. The benefits are the following:

- i) it allows for an analysis of the most common macro level shocks using data that is founded on micro-level firm-specific investment decisions,
- ii) unlike all national accounts based micro-level data for the UK, it allows for using a quarterly frequency, which is a pre-requisite for the analysis on the effects of shocks, and
- iii) it potentially allows for more detailed splits of the dataset than conventional macro-level investment data in terms of sector, firm size and the effects of other, for example corporate finance related, variables on investment.

The micro data used in the study is firm-level data on financial account items, sourced from the proprietary Standard & Poor's (S&P) Capital IQ database.⁶ The database includes over 300,000 UK firms, and has financial account data available for a subsample of these. For the vast majority of the firms in the database, no financial account data is available, and therefore they cannot be used for the analysis. Indeed, for the purposes of the analysis in the current study, a sample of about 3,000 firms was selected based on the availability of relevant financial account data for these firms. While this may seem a relatively small sample, one needs to keep in mind that the ONS data on business investment is based on a sample of around 27,000 firms (ONS (2013)), and in fact, the capital expenditure carried out by the firms in the current study was around 75% of UK business investment in 2014.⁷

For the comparison of micro and macro-level data, a decision needs to be made on the exact firm-level financial account variables that most closely correspond to the macro level variables (which are based on ONS data). Obviously, a one-to-one matching cannot be achieved, as the micro and macro level variables measure different things. However, the decisions on the variables follow fairly standard conventions used in previous micro-data literature, while taking into account data availability issues in the Capital IQ database. Ultimately, the test of the relevance of the chosen financial account variables is their correlation with the corresponding macro level variables (see next section).

Table 2 details the specific financial account variables used in the study as proxies for the corresponding macro level data. It also includes weights for the sectoral splits used in the analysis. In terms of the chosen micro variables, operating income (from the income statement) is used as a proxy for GDP, as it most closely matches gross value added on the firm level. For investment, capital expenditure (from the cash flow statement) is used, which is a fairly obvious match. For the cost variable, total operating expenses (income account) includes all relevant costs in generating the firm's income, including costs of goods sold and wages. Hence, an average of CPI and wages (measured by average weekly earnings) are used as the corresponding macro variable. All the firm-level data is in

⁶ Disclaimer of Liability Notice for use of the database: The database may contain information obtained from third parties, including ratings from credit ratings agencies such as Standard & Poor's. Reproduction and distribution of third party content in any form is prohibited except with the prior written permission of the related third party. Third party content providers do not guarantee the accuracy, completeness, timeliness or availability of any information, including ratings, and are not responsible for any errors or omissions (negligent or otherwise), regardless of the cause, or for the results obtained from the use of such content. THIRD PARTY CONTENT PROVIDERS GIVE NO EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE. THIRD PARTY CONTENT PROVIDERS SHALL NOT BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, EXEMPLARY, COMPENSATORY, PUNITIVE, SPECIAL OR CONSEQUENTIAL DAMAGES, COSTS, EXPENSES, LEGAL FEES, OR LOSSES (INCLUDING LOST INCOME OR PROFITS AND OPPORTUNITY COSTS OR LOSSES CAUSED BY NEGLIGENCE) IN CONNECTION WITH ANY USE OF THEIR CONTENT, INCLUDING RATINGS. Credit ratings are statements of opinions and are not statements of fact or recommendations to purchase, hold or sell securities. They do not address the suitability of securities or the suitability of securities for investment purposes, and should not be relied on as investment advice.

⁷ A caveat to this is that some proportion of investment by the firms used in the analysis could have taken place abroad.

nominal terms, and the filtered proxy variables are transformed into real terms by deflating them with the GDP deflator (apart from the cost series, which attempts to take into account changes in prices).

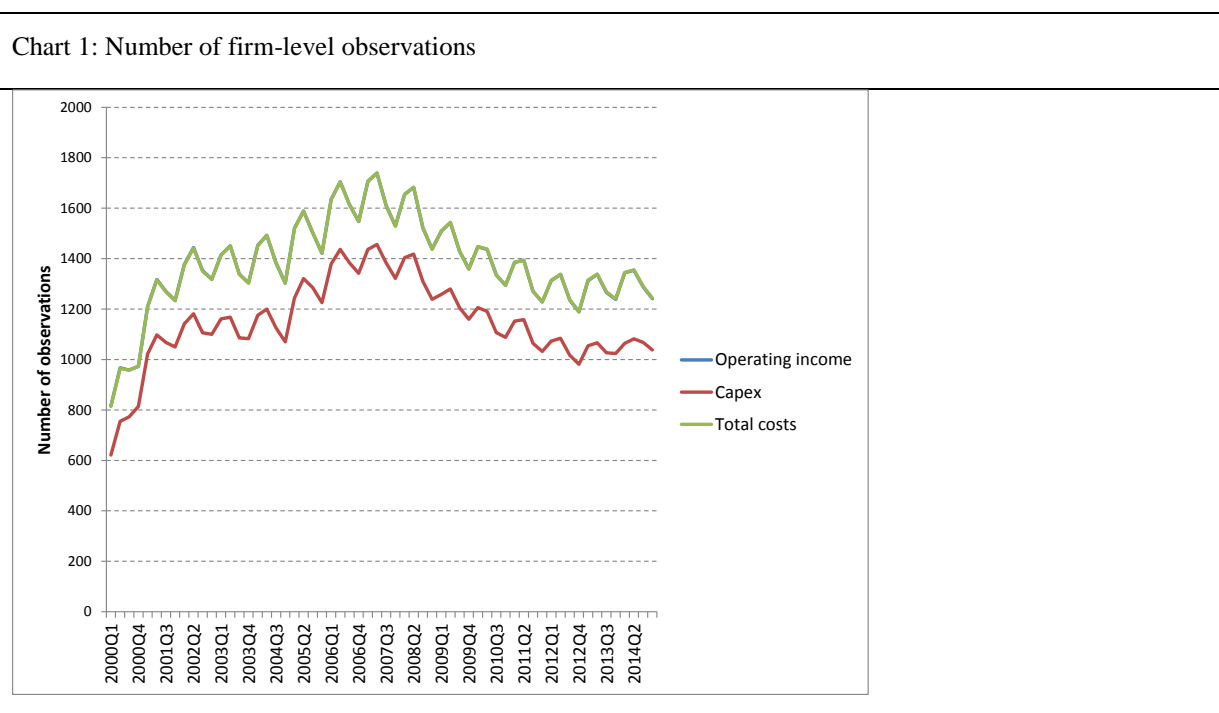
Table 2: Variable and sector definitions

Link between micro and macro variables:				
<u>Capital IQ variables (abbreviation):</u>		<u>Macro variables (from ONS):</u>		
Operating income (Ropinc)		Real GDP		
Capital expenditure (Rcapex)		Real business investment		
Total operating expenses (Costs)		Cost price index (average of CPI + average weekly earnings)		
S&P sector classifications and weights (based on total revenue in 2014):				
	Weight:	Sub-sector used in analysis:	Abbrvtn:	Weight:
Financials	0.13	Financials	FIN	0.13
Consumer Staples	0.15	Consumer goods and services	CON	0.52
Consumer Discretionary	0.20			
Healthcare	0.03			
Industrials	0.13			
Energy	0.16			
Utilities	0.06	Industrial	IND	0.29
Materials	0.07			
Information Technology	0.02	IT goods and services	IT	0.06
Telecommunication Services	0.04			
		Large firms (>250 employees)	BIG	0.96
		SMEs (<251 employees)	SME	0.04

One of the advantages of micro data is that it can be used to analyse smaller sub-samples of the data, which can be interesting if firms respond to shocks in a heterogeneous way. In the sectoral analysis carried out here, the firms are divided into six partly overlapping categories (called sub-sectors from now on). First, the firms are divided into four broad industrial categories and second, into two categories based on their size (determined on the basis of the total number of employees reported by the firm in 2014). The split in size conforms to the usual definition of SMEs (less than 250 employees), while the industrial categories are an attempt to group the original S&P categories into uniform “end-user” type groups.

Due to data limitations, it is not possible to achieve an exact correspondence with the definitions of national account sectors, and one of the sectors (consumer goods and services) ends up accounting for

more than half of the total revenue of all the firms in the sample. Nevertheless, the weights of the firms in different sectors are roughly similar to those found in national accounts (NA) data; service sector firms account for around 60-65% of the sample, compared to around 75% in NA data. While the weight of SMEs is clearly smaller than in NA data (around 5% in the sample versus around 45% in NA data), the number of SMEs in the dataset (around 500 firms) is sufficiently large for inference to be meaningful. Overall, even if the characteristics of the sample do not fully correspond with NA data, the sub-sample split allows for analytically interesting heterogeneity in the sample and is justified on the basis of attempting to explain differences in “end-user” behaviour (rather than “production” behaviour in NA data).⁸



The actual number of firms included in the analysis over time for the different variables is presented in Chart 1. The numbers are very similar for total costs and operating income (overlapping each other in the Chart), both slightly higher than for capital expenditure. The numbers are also fairly constant after a sizeable increase during the first year of the sample. This also forces a natural cut on the sample period, as the number of firms for which the data is available is dramatically lower before the year 2000.

⁸ In other words, the analysis is able to answer a question like “what are the effects on investment of shock x in a consumer-facing sector?” rather than “what are the effects on investment of shock x in a sector producing type z consumer goods?”.

In the VAR framework I am using, all the variables are derived as the filtered factors as described above, apart from the interest rate variable. It can be viewed as the “observable” variable in a FAVAR type framework. However, selecting the interest rate variable to be used in the models is not straightforward,⁹ especially given the firm-level environment in which the model operates. It is fairly obvious that traditional models where the monetary policy rate is included as a direct measure of monetary policy shocks, will fail to account for the unusual dynamics of monetary policy and business investment in the post-financial crisis environment. This is partly because unconventional measures that have been introduced in recent years do not show up in the monetary policy rate, but also because of the surprisingly weak estimated reaction of investment to the very low policy rate in many advanced economies.

To examine the effects of using different interest rate variables, three candidates better suited to the current monetary policy environment are used:

- i) An aggregate measure, calculated in-house at the Bank of England, of the real cost of capital (RCC) is used in the benchmark model.¹⁰ This measure is constructed from the weighted average cost of capital (WACC), relative investment prices, depreciation and a tax adjustment factor.¹¹
- ii) For some of the robustness analysis, an in-house Bank of England measure of the Bank Rate that takes into account the effects of unconventional monetary policy measures since the financial crisis is also used (shadow R). Shadow R is based on Bank estimates of the effects of unconventional monetary policy measures (see Bridges and Thomas (2012) and Joyce et al. (2011)). While no mapping of unconventional measures onto policy rate space is perfect, it is likely to be a better measure of monetary policy stance than the pure Bank Rate, and hence, is a viable option for robustness analysis.
- iii) Another important factor for investment dynamics in recent years has been the role of large fluctuations of macroeconomic uncertainty. There is abundant evidence in the literature that uncertainty affects firms’ investment decisions, because it increases the option value of waiting before committing to a new investment project (see, for example, Bernanke (1983), Bloom and Van Reenen (2007), Bloom (2009) and Jurado et al. (2015)). I take into account this uncertainty with a proxy for aggregate as well as sub-sector level uncertainty-adjusted cost of capital by combining the cost of capital with a stock-price volatility based uncertainty measure introduced by Gilchrist et al. (2013) (for details of the measure, see Appendix 1). The combination is done by taking the first

⁹ There is a long-standing literature on the effects of monetary policy rates on business investment. Pioneering studies by Bernanke and Gertler (1995) and Christiano et al. (1999) found a negative effect of investment on positive policy shocks for the US, while other studies have discussed the lack of a negative relationship between investment and cost of capital (see Abel and Blanchard (1986), Gilchrist and Zakrajsek (2007) and Schaller (2006)).

¹⁰ Ideally, a firm-level cost of capital measure should be used in the models. However, as no such universally used measure exists, this is left for future research.

¹¹ See Appendix 1 for details on the construction of RCC.

principal component of the aggregate RCC and 7 different uncertainty measures (1 aggregate, 6 sub-sector measures). This gives a relatively neutral way of estimating a sector-specific cost of capital measure that takes into account the marginal contribution of uncertainty in addition to RCC,¹² while also focusing on “pure” uncertainty rather than volatility.¹³ (See different RPC series in Appendix 3, Chart h).

4. Results

This section presents the results of the benchmark analysis as well as highlights some interesting extensions. For the factor models on the firm-level data (see charts a)-c) in Appendix 3), it is worth highlighting that the resulting factors are strongly positively correlated, with correlation coefficients of around 0.7, with the corresponding macro series. As detailed above, there are a number of advantages to using firm-level data, but nevertheless, it is reassuring to see that a positive correlation with the macro series exists. This is a crucial pre-condition for the viability of the analysis, which then allows me to also examine interactions between variables and cuts of the micro-level data unavailable when using aggregate level national accounts data.

In the following subsections, I turn to the results from the VAR analysis.

4.1 Benchmark VAR results

The impulse responses from the VAR model estimation are shown in Charts 2 to 4. To a large extent, these exhibit dynamics that are in line with previous literature as well as a macro-variable based model of the UK economy for the same time period. In general, supply and cost of capital (RCC from now on) shocks have a more persistent effect on investment than demand shocks. The impulse responses are relatively similar across the six sub-sectors considered in the analysis, but there are more differences in the forecast error variance decompositions (FEVD) shown in Chart 4.

In terms of the sub-sector results, the largest sub-sector, consumer goods and services, does not particularly stand out from the other sectors. Based on both the IRF and FEVD analysis, supply shocks tend to be more relevant (i.e., they have had larger and more persistent effects) in explaining investment dynamics than the other three shocks for the consumer sector.

¹² While no measure is perfect, this method offers a simple and transparent way of taking uncertainty into account in the modelling framework. I experimented with different methods of combining the RCC and uncertainty measures, but there is no clear advantage to moving to more complex methods, as differences between the results are typically small, and adding complexity appears difficult to justify on either theoretical or practical grounds. The first principal component explains about 80-90% of the common variation of the RCC and the uncertainty measure, which also suggests that this simple method is appropriate.

¹³ There is an important distinction made in recent literature between uncertainty and volatility (or risk); the former is unforecastable whereas the latter may not be (see e.g. Jurado et al. (2015) and Orlik and Veldkamp (2014)). The measure I use is purged from the forecastable (market-returns based) component and can hence be seen as a measure of uncertainty rather than volatility. Nevertheless, the measure should be seen as illustrative rather than definitive, as further advances into uncertainty literature are beyond the scope of the current study.

Two sectors that stand out as having distinctively different innovation accounting results are the IT and financial sectors. For the IT sector, the effects of supply as well as RCC shocks are more persistent than for other sectors, and in the FEVD analysis, RCC shocks account for a large share of investment variation in the IT sector. The latter result could be related to the fact that indebtedness of IT sector firms has tended to be relatively high compared to the other sectors¹⁴ and hence, investment decisions in IT firms have probably been more sensitive to changes in the RCC. In contrast, demand shocks tend to have the largest initial effect, as well as account for more of the investment variation for the financial sector, which could suggest it is more exposed to cyclical fluctuations than the other sectors. In terms of the FEVD analysis, supply shocks are much less relevant for financial firms than demand shocks.

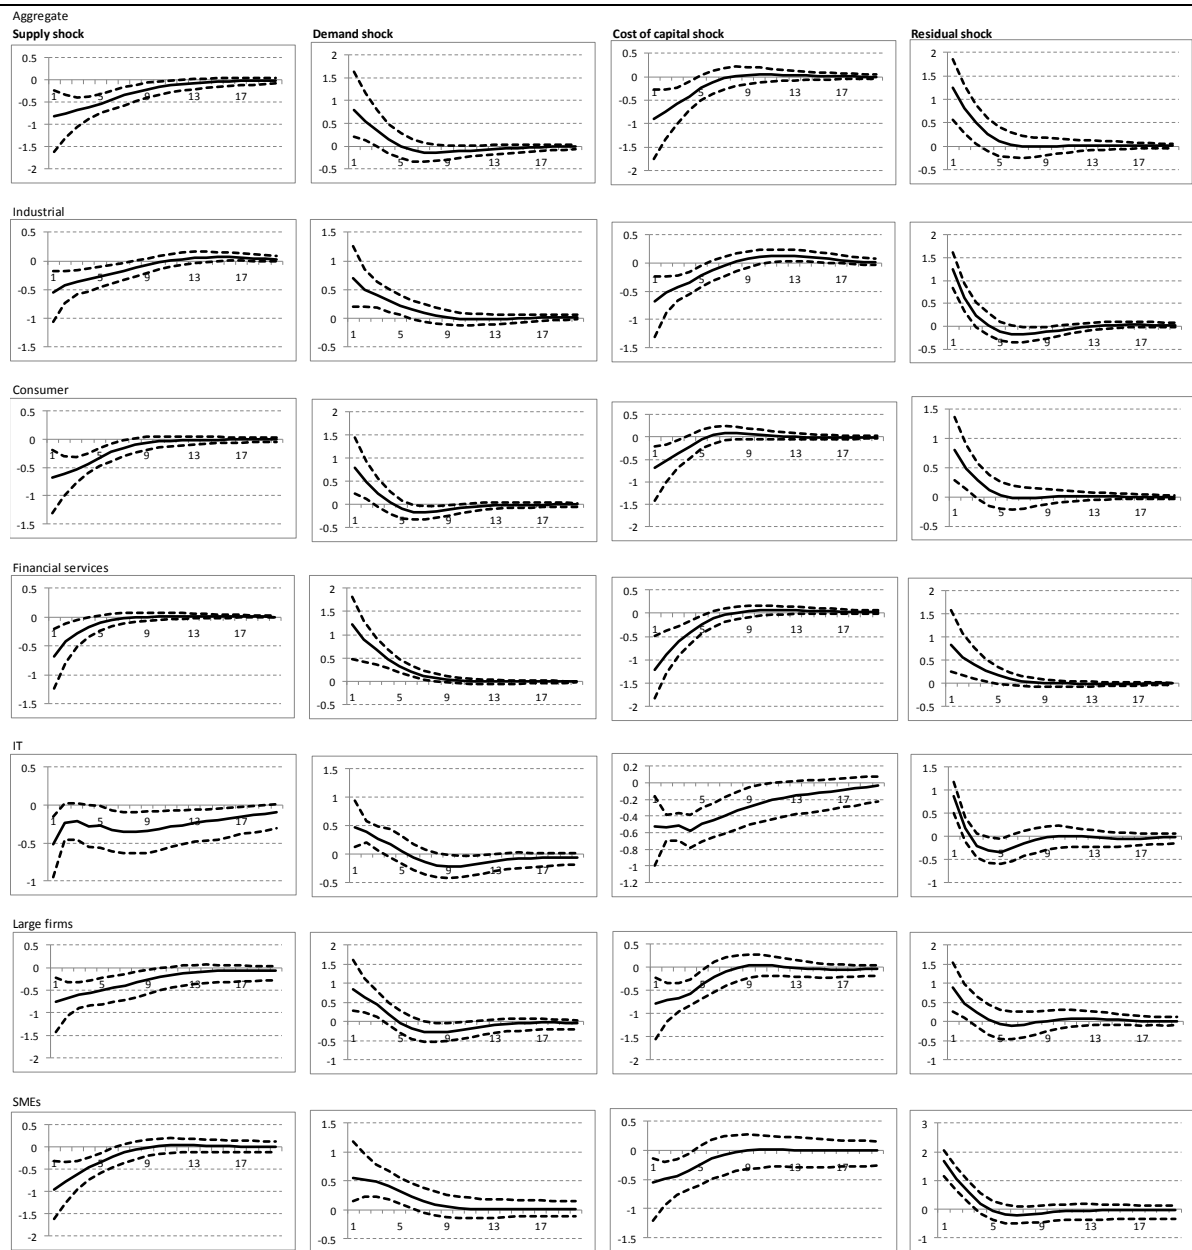
There are also some differences in the results between large firms and SMEs. The effects of supply shocks are more persistent for large than small firms, and the FEVD analysis suggests that supply shocks account for a much larger share of investment variation for the large firms than SMEs. Demand and RCC shocks also account for a larger share of investment variation for the large firms than SMEs. The lower share of RCC shock variability could be related to the fact that indebtedness of the SME sector has been considerably lower than that of the larger firms over the sample period.¹⁵ On the other hand, the residual shock accounts for a larger share of investment variation for SMEs than for large firms (or any of the other sub-sectors). This could point to the more heterogeneous nature of smaller firms, whose investment decisions are possibly more affected by idiosyncratic factors not captured by the model framework.

For an illustrative comparison, results with a Choleski ordering for the aggregate model are presented in Appendix 5. Signs of the impulse responses are theoretically correct. The height of the effects of the shocks on investment comes through later than in the sign-restricted results, partly due to the fact that investment cannot react to the shocks during the same quarter. However, both methods suggest that the effects of supply and cost shocks are more persistent than the other shocks.

¹⁴ As measured by the debt/equity ratio of listed UK firms, which was around 2.5 on average for IT firms compared to around 1 for other firms during the sample period.

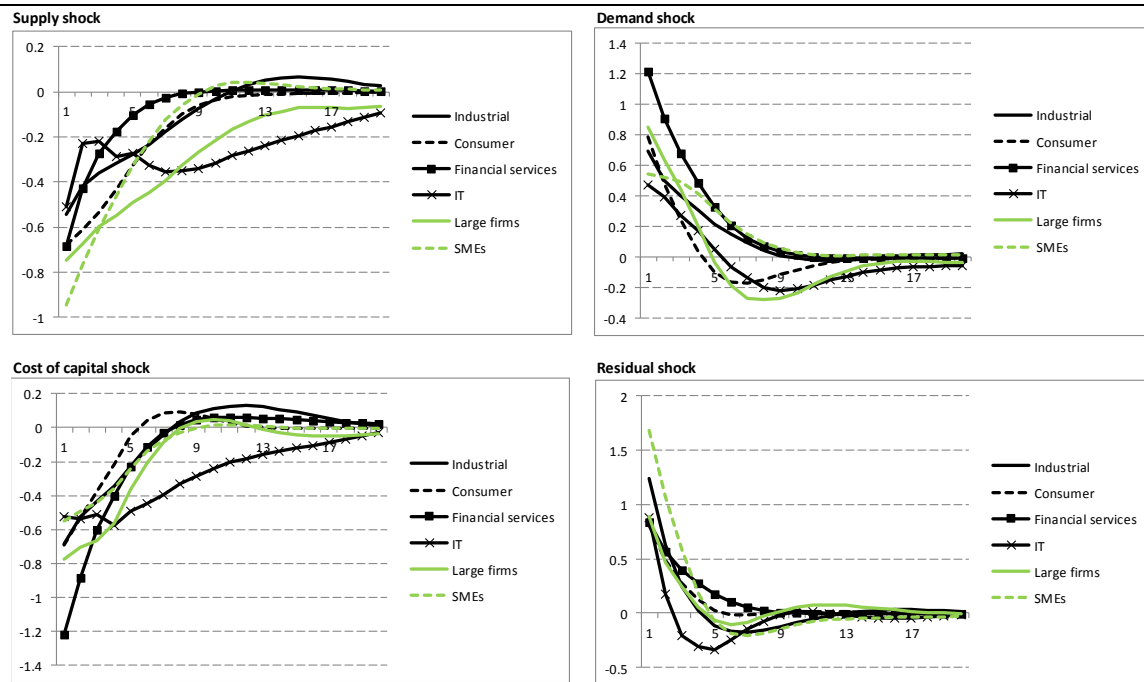
¹⁵ The debt/equity ratio of large firms was around 1.1 on average and SMEs around 0.6 during the sample period.

Chart 2: Impulse responses of investment with 68% confidence intervals



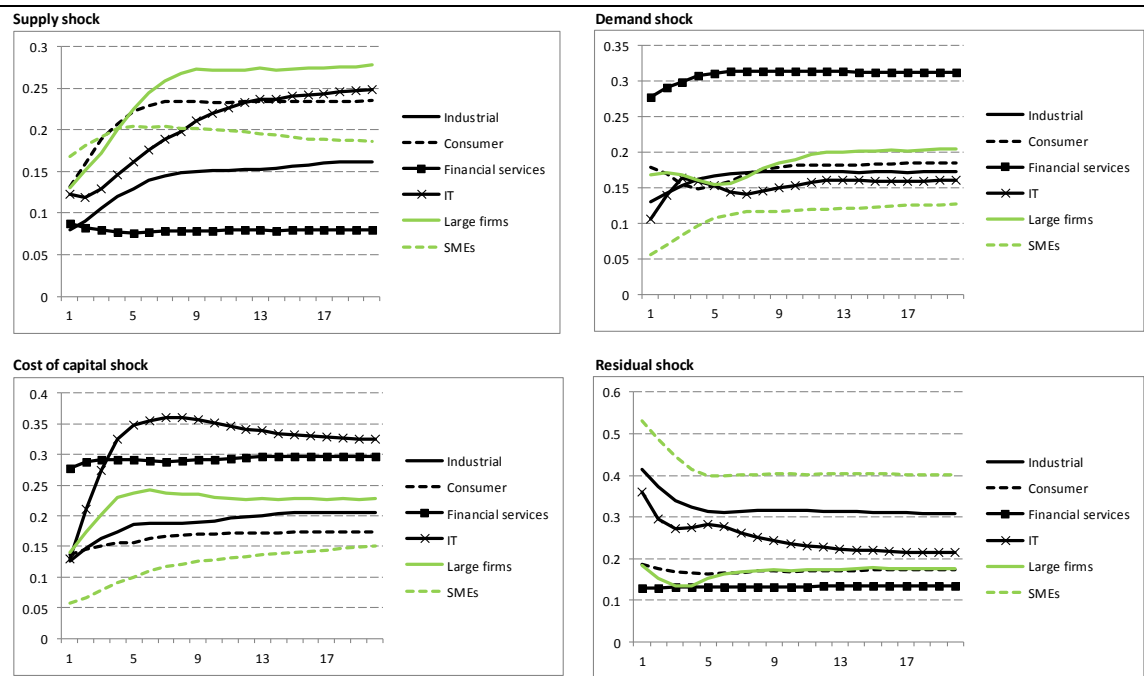
The charts show change in the (log) level of investment for 1-standard deviation shocks over 20 quarters.

Chart 3: Impulse responses – comparison across sectors and size of firm



The charts show change in the (log) level of investment for 1-standard deviation shocks over 20 quarters.

Chart 4: Forecast error variance decompositions – comparison across sectors and size of firm



The charts show median contributions to the variance of investment for 1-standard deviation shocks over 20 quarters. Given the nature of the Bayesian estimation method for the posterior distributions, the decompositions do not sum to one across the shocks.

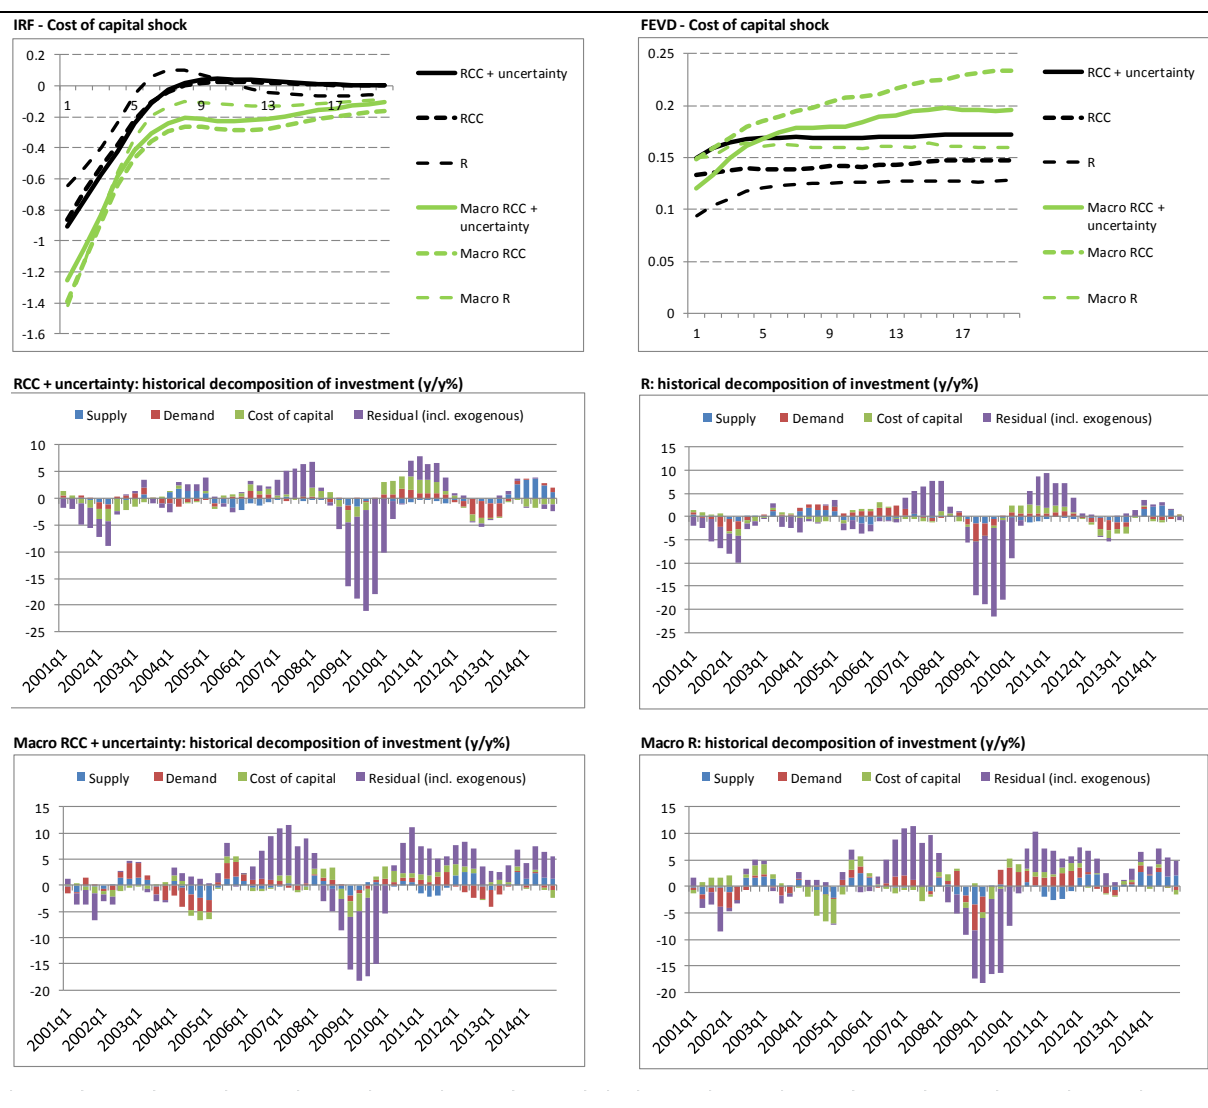
4.2 Interest rate shocks

Given the uncertainty related to the interest rate variable used in the model, it is also informative to examine the results related to the RCC shock and their robustness to the choice of the variable. Chart 5 shows the IRFs, FEVDs and historical decompositions of both the micro-data based model and a corresponding macro model for three different interest rate variables; i) a pure monetary policy rate (shadow R, which also takes into account quantitative easing measures introduced after the financial crisis), ii) an RCC cost of capital interest rate and iii) RCC with an added uncertainty variable (as detailed above). The differences in the IRFs are not large for the different interest rate variables, but the FEVD results are more diverse. In particular, for the micro models, it is interesting to note that adding more information (in other words, moving from R to RCC to RCC + uncertainty) increases the share of variance of investment explained by the interest rate variable.¹⁶ For the macro models – which are not the main focus for this study and should be taken as indicative only – the model with the R variable also has the lowest explanatory power for the variance in investment accounted for by the interest rate variable. Overall, these results suggest, in line with some of the previous literature, that a traditional model with a monetary policy rate attributes a larger proportion of investment dynamics to residual shocks and hence, is less appropriate for explaining recent dynamics in investment in the UK than models that truly attempt to account for the cost of capital faced by firms.

Finally, the historical decompositions suggest that the lower level of cost of capital explained most of the rebound in investment growth after the financial crisis, but this support has since faded. The dissipation of the spike in uncertainty caused by the financial crisis was probably a significant factor in the rebound in investment. However, the negative demand shock caused by the euro area crisis contributed negatively to investment growth in 2012-2013. Positive supply shocks, which could have been related to an improvement in bank lending conditions and some easing in persistent resource reallocation constraints caused by the financial crisis, supported investment towards the end of the sample.

¹⁶ The results are also consistent with those of Bloom et al. (2007) in terms of including the uncertainty indicator in the model reduces the responsiveness of investment to demand shocks.

Chart 5: Comparisons of shock effects and decompositions with different cost of capital variables



4.3 Model extensions

To highlight the flexibility of the modelling framework, this section reports the results of two relevant extensions to the benchmark model introduced above. First, I examine the effects of corporate payout shocks and second, I illustrate some results from a time-varying version of the model.


4.3.1 Corporate governance shocks

The framework can also be used to answer other questions related to firm-level drivers of investment decisions. For example, one can examine the effects of corporate dividend policy on investment. There is a long-standing debate in the literature on these effects; while in theory, dividends and investment should not be substitutes according to Modigliani-Miller type corporate behaviour (see, for example, Kliman and Williams (2014)), some authors have found evidence of higher dividend

payments crowding out investment (see, for example, Blundell-Wignall and Roulet (2014). Other authors also stress the heterogeneity of behaviour across firms (Mathur et al. (2015)).

For this analysis, an aggregate filtered indicator of firm-level payouts (which is defined as dividends plus equity buy-backs, as is customary in the literature) was constructed in a similar fashion to the other aggregate indicators in the model. This indicator is shown (alongside macro business investment) in panel d) of Appendix 3. There was an acceleration in the growth rate of real payouts before the financial crisis, a subsequent collapse, and then a recovery towards the end of the sample. This is in line with the widely reported narrative of payout dynamics over the past decade.

The IRFs¹⁷ of the model that includes the payout as a 5th variable are shown in Chart 6. These results do not support the crowding out effect - even though for the sign-restricted model the response of investment is initially negative, it is not statistically significant. With the recursive ordering, there is no clear statistically significant response either. The results (not shown) across the sub-sectors are similar to the aggregate result. Hence, based on the firm-level VARs, it does not appear to be the case that dividends and share buybacks have crowded out investment in the UK over the past 15 years.



4.3.2 Time-varying VAR analysis

Given the potential for changes in investment dynamics since the financial crisis period, an analysis of the aggregate level model was also carried out within a sign-restricted time-varying VAR (TVAR)

¹⁷ For the sign-restricted model, no restrictions on the responses of the other variables are added for the payout shock, as theoretically the direction is not clear. Intuitively, payouts are restricted to respond positively to a demand shock and negatively to a supply shock, given that more (the former case)/less (the latter case) funds are available to distribute to shareholders.

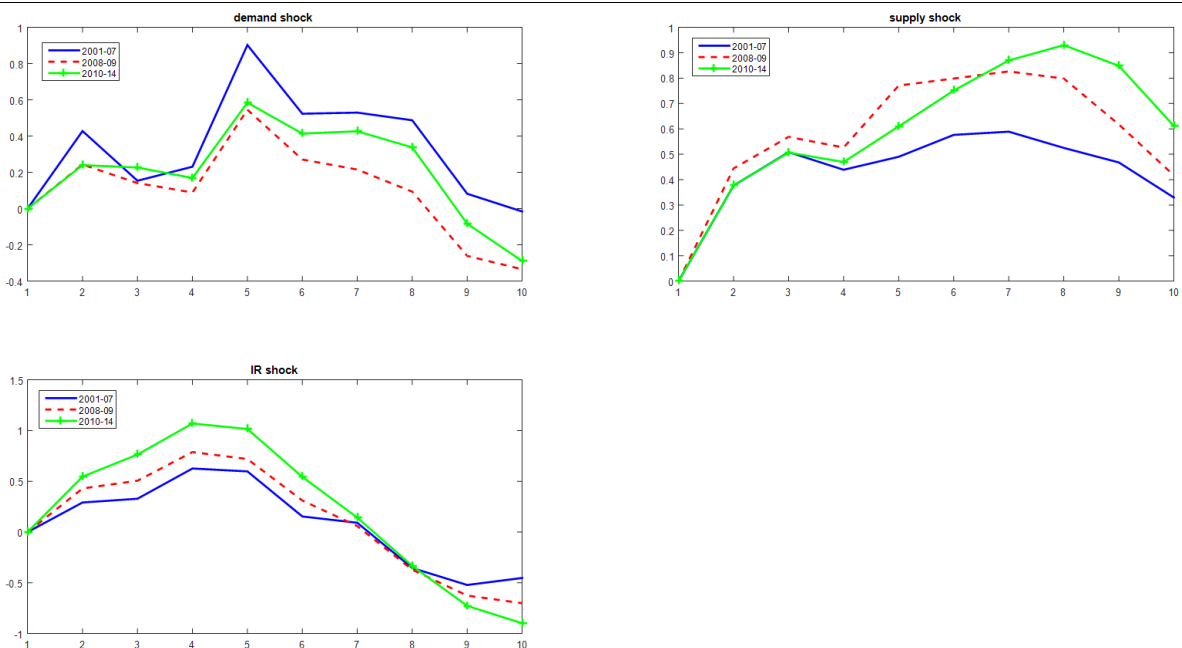
framework. For this purpose, a recent non-parametric technique developed by Giraitis et al. (2014) is used. This technique uses a kernel-based method for estimating random time-varying coefficients, rather than using more conventional state space models. The authors show that this method has the advantage of yielding parameter estimates that are consistent and asymptotically normal, while allowing for smoothly changing estimates.

The TVAR model is estimated using the same basic features, including the same sign restrictions, as the benchmark model. However, given that the sample in the model is relatively short, some of the TVAR results are volatile and not always robust to changes in the technical specifications. Hence, the results presented here are best viewed as illustrative¹⁸. In Chart 7, the IRFs and FEVDs are divided into three main periods; pre-crisis (2001-2007), crisis (2008-2009) and post-crisis (2010-2014), with the means of the relevant responses for each sub-sample presented in the chart. There is a tendency for demand shocks to be more persistent, and account for a larger share of variation in investment in the pre-crisis period, whereas RCC and supply shocks have become more relevant since then. The latter result may be indicative of more persistent effects of both negative and positive supply shocks on investment (and ultimately productivity potential) since the financial crisis. Furthermore, it is interesting to note that the larger post-crisis responsiveness to interest rate shocks is only present in the benchmark version of the model including the uncertainty augmented cost of capital, whereas with a pure RCC shock, the results are more in line with pre-crisis ones. Hence, it could be that investment has been more sensitive to increased firm-level uncertainty in recent years.

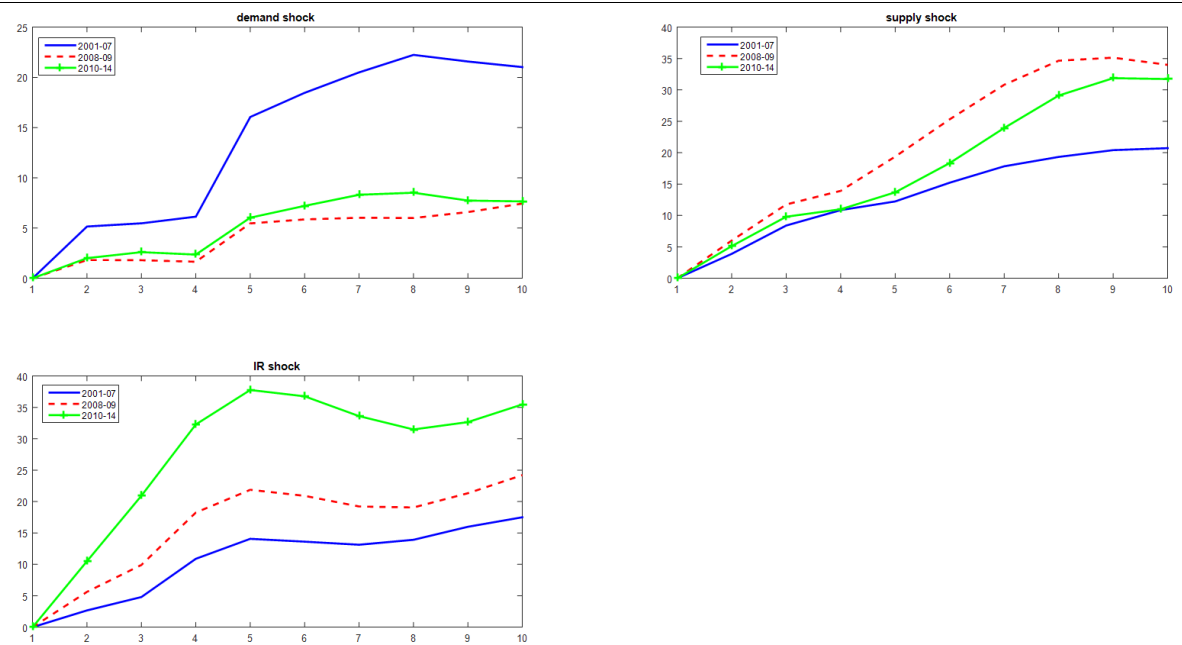
¹⁸ The main choice in the estimation concerns the kernel bandwidth. For the results shown here, a bandwidth of 0.8 was chosen, although they are relatively robust to changes between 0.7 and 0.9 (as expected for the sample size).

Chart 7: Impulse responses (IRF) and forecast error variance decompositions (FEVD) in the TVAR

a) Investment IRFs



b) Investment FEVDs



The results are from a model with a bandwidth of 0.8. The charts show the average effects over the time period indicated (with no effect in impact quarter 1) in the charts for 1-standard deviation shocks over 10 quarters. Residual shock is not reported, and hence FEVDs in the chart do not sum to one.

5. Conclusions

This paper analyses drivers of business investment in the UK in a simple VAR framework by filtering firm-level micro data into macro-level indicators. A traditional innovation accounting analysis is then carried out on the VAR models that include the indicators, also allowing for cuts on the data based on firm sectors, size, cost of capital and time.

The analysis brings up several interesting results, some of which are worth highlighting at a more general, policy-relevant level. Supply shocks are estimated to have been more persistent and more important in explaining investment dynamics in the UK than demand shocks over the past 15 years, and their importance appears to have increased since the financial crisis. Furthermore, it is also worth noting that shocks to the cost of capital, and uncertainties related to it, have generally been more important for firms in sectors with higher indebtedness, whereas corporate governance issues as measured by dividend payments and share buybacks do not appear to have been a major driver of investment. Long-term policies that foster a favourable environment for investment while minimising uncertainty appear best-placed to lead to an improvement in investment and ultimately potential productivity growth.

There are a number of avenues for future research, both in terms of interpreting recent weakness in investment dynamics in the UK as well as more globally. First, as suggested by the results of the current study, firm-level heterogeneity can matter and hence, there is clearly room for more detailed firm-level analysis. In particular, examining the role of the cost of capital, expected returns and firm-level uncertainty in investment decisions appears a particularly important topic. Furthermore, investigating the links between investment and productivity with firm-level data for the UK economy is another vital area of research, especially given the shocks that have hit the economy over the past decade.

References

Abel, A. and O. Blanchard (1986), “The Present Value of Profits and Cyclical Movements in Investment”, *Econometrica*, 54, pp. 249-273

Baumann, U. and S. Price (2007), “Understanding investment better: insights from recent research”, *Bank of England Quarterly Bulletin Article 2007Q2*

Bernanke, B. (1983), “Irreversibility, Uncertainty, and Cyclical Investment”, *Quarterly Journal of Economics*, 98, pp. 85-106

Bernanke, B., J. Boivin and P. Elias (2005), “Measuring the effects of monetary policy: a factor-augmented vector autoregressive (FAVAR) approach”, *The Quarterly Journal of Economics*, February 2005, pp. 387-422

Bernanke, B. and M. Gertler (1995), “Inside the black box: the credit channel of monetary policy transmission”, *Journal of Economic Perspectives* 9(4), pp. 27-48

Bloom, N. (2009), “The Impact of Uncertainty Shocks”, *Econometrica*, 77, pp. 623-685

Bloom, N., S. Bond and J. Van Reenen (2007), ‘Uncertainty and investment dynamics’, *Review of Economic Studies*, Vol. 74, pp. 391–415

Blundell-Wignall, A. and C. Roulet (2015), “Infrastructure versus other investments in the global economy and stagnation hypotheses: What do company data tell us?”, *OECD Journal: Financial Market Trends*, Vol. 2014/2

Bond, S., Elston, J.A., Mairesse, J. and B. Mulkey (2003), “Financial factors and investment in Belgium, France, Germany, and the United Kingdom: a comparison using company panel data”, *The Review of Economics and Statistics*, 85(1), pp. 153–165

Bond, S., A. Klemm, R. Newton-Smith, M. Syed and G. Vlieghe (2004), “The roles of expected profitability, Tobin’s Q and cash flow in econometric models of company investment”, *Bank of England Working Paper No. 222*

Bridges, J. and R. Thomas (2012); "The impact of QE on the UK economy – some supportive monetarist arithmetic", Bank of England Working Paper No. 442

Canova, F., and M. Paustian (2011), "Business cycle measurement with some theory", *Journal of Monetary Economics*, 58, pp. 345-361

Christiano, L., M. Eichenbaum and C. Evans (1996), "The Effects of Monetary Policy Shocks: Evidence from the Flow of Funds", *The Review of Economics and Statistics*, 78(1), pp. 16-34

Christiano, L., M. Eichenbaum, and C. Evans (1999); "Monetary policy shocks: What have we learned and to what end?" in *Handbook of Macroeconomics*, ed. by J. B. Taylor and M. Woodford, Elsevier, vol. 1 of *Handbook of Macroeconomics*, chapter 2, pp. 65-148

Doz, C., D. Giannone and L. Reichlin (2006), "A two-step estimator for large approximate dynamic factor models based on Kalman Filtering", Unpublished manuscript, Université Libre de Bruxelles

Duchi, F. and A. Elbourne (2016), "Credit supply shocks in the Netherlands", *Journal of Macroeconomics*, 50, pp. 51-71

Farrant, K., M. Inkinen, M. Rutkowska and K. Theodoridis (2013), "What can company data tell us about financing and investment decisions?", Bank of England Quarterly Bulletin Article 2013Q4

Giannone, D., Reichlin, L. and D. Small (2008), "Nowcasting: The real-time information content of macroeconomic data", *Journal of Monetary Economics*, 55, pp. 665-676

Gertler, M. and P. Karadi (2011), "A model of unconventional monetary policy", 58(1), pp. 17-34

Gilchrist, S., J. Sim and E. Zakrajsek (2014), "Uncertainty, Financial Frictions, and Irreversible Investment", NBER Working Paper No. 20038

Gilchrist, S. and E. Zakrajsek (2007), "Investment and the cost of capital: new evidence from the corporate bond market", NBER Working Paper No. 13174

Giraitis, L., G. Kapetanios and T. Yates (2014), "Inference on stochastic time-varying coefficient models", *Journal of Econometrics*, 179, pp. 46-65

Hall, R. and D. Jorgenson (1969), “Tax policy and investment behaviour”, *American Economic Review*, 57, pp. 391–414

Inkinen, M., Stringa, M. and K. Voutsinou (2010), “Interpreting equity price movements since the start of the financial crisis”, *Bank of England Quarterly Bulletin Article 2010Q1*

Jorgenson, D. (1963), “Capital theory and investment behaviour”, *American Economic Review*, 53, pp. 247–56.

Joyce, M., M. Tong, and R. Woods (2011); “The United Kingdom’s quantitative easing policy: design, operation and impact”, *Bank of England Quarterly Bulletin 2011Q3*, pp. 200-212

Jurado, K., Ludvigson, S. and S. Ng (2015), “Measuring Uncertainty”, *American Economic Review*, 105(3), pp. 1177-1216

Kliman, A. and S.D. Williams (2014), “Why ‘financialisation’ hasn’t depressed US productive investment”, *Cambridge Journal of Economics*, 39, pp. 67-92

Litterman, R. (1986), “Forecasting with Bayesian Vector Autoregressions – Five years of experience”, *Journal of Business and Economic Statistics*, 4 (1986), pp. 25-38

Lutkepohl, H. (2005), “New introduction to multiple time series analysis”, Springer-Verlag

Mathur, A, N. Rao, M. Strain and S. Veuger (2015), “Dividends and Investment: Evidence of Heterogeneous Firm Behavior”, NYU Wagner Research Paper No. 2575860

Melolinna, M. (2012), “Macroeconomic shocks in an oil market VAR”, *ECB Working Paper No. 1432 (May 2012)*

Mountford, A. and H. Uhlig (2009), “What are the effects of fiscal policy shocks?”, *Journal of Applied Econometrics*, 24, pp. 960-992

ONS (2013), “Business Investment Quality and Methodology Information”, Office for National Statistics Information Paper, 19th April 2013

Orlik, A. and L. Veldkamp (2014), “Understanding uncertainty shocks and the role of black swans”, NBER Working Paper No. 20445

Peersman, G. (2005), “What Caused the Early Millennium Slowdown? Evidence Based on Vector Autoregressions”, *Journal of Applied Econometrics*, 20, pp 185-207

Peersman, G. and F. Smets (2001), “The monetary transmission mechanism in the euro area: more evidence from VAR analysis”, ECB Working Paper No. 91

Rubio-Ramirez, J., D. Waggoner and T. Zha (2010), “Structural Vector Autoregressions: Theory of identification and algorithms for inference”, *The Review of Economic Studies*, 77, pp. 665-696

Schaller, H. (2006), “Estimating the Long-Run User Cost Elasticity”, *Journal of Monetary Economics*, 101, pp. 725-736

Uhlig, H. (2005), “What are the effects of monetary policy on output? Results from an agnostic identification procedure”, *Journal of Monetary Economics*, 52, pp. 381-419

Appendix 1: Cost of capital and uncertainty indicators

Real cost of capital

The real cost of capital (RCC) measure is based on the following formula (see Hall and Jorgenson (1969) and Schaller (2006)):

$$RCC = P * T * (WACC - d) \quad (A1)$$

where P is the relative price of investment goods (investment price deflator divided by GDP deflator), T is a tax adjustment factor (taxes minus allowances, as detailed in Schaller (2006), using the main corporate tax rates and investment allowances for plant, machinery and equipment in the UK), WACC is weighted average cost of capital (see below) and d is depreciation (assumed to be 6% p.a., based on historical estimates for the UK).

The WACC measure (in real terms) is given by the following formula:

$$WACC = [(r^f + ERP) * w^{equity} + (r_{loan}^{debt} - \pi) * w_{loan}^{debt} + (r_{bond}^{debt} - \pi) * w_{bond}^{debt}] \quad (A2)$$

The cost of equity is composed of the risk free rate (r^f) (10-year real rate from inflation swaps) plus an equity risk premium (ERP), which is calculated using an in-house Dividend Discount Model (DDM). The DDM also takes into account share buybacks, time-varying long-run growth estimates and the entire yield curve (for a description on how the DDM works, see Inkinen et al. (2010)). The equity risk premium is calculated for companies publicly listed in the UK. The nominal cost of bank loan debt, r_{loan}^{debt} , is calculated as the average debt interest expense divided by the book value of total bank loan debt on the company's balance sheet and for market-based bond debt, yields on sterling BBB-rated bonds of UK non-financial corporations with maturity of 8-12 years (r_{bond}^{debt}). The real cost of debt is then derived by subtracting inflation rate π , proxied by 10-year inflation rate from inflation swaps. The weights are total shareholders' equity (w^{equity}), bank loan debt (w_{loan}^{debt}) and bond debt (w_{bond}^{debt}) as shares of total assets.

Stock-price based firm-level uncertainty indicator¹⁹

The estimate of firm-level uncertainty is based on a two-step procedure. First, the forecastable variation in daily firm-specific excess returns over a market portfolio is calculated as follows:

$$(R_{i,t,d} - r_{td}^f) = \alpha_i + \beta_i(R_{M,t,d} - r_{td}^f) + u_{i,t,d} \quad (A3)$$

¹⁹ This section draws heavily on Gilchrist et al. (2014).



where $R_{i,td}$ is the daily stock price return for firm i in day d during quarter t , $R_{M,td}$ is the corresponding market (in this case, FTSE350) return, r_{td}^f is the risk-free short-term interest rate (3-month LIBOR), α_i is the firm-specific alpha, β_i is the firm-specific beta and $u_{i,td}$ is the daily idiosyncratic return.

In the second step, a quarterly measure of firm-specific standard deviation σ_{it} is calculated from the daily idiosyncratic returns:

$$\sigma_{it} = \sqrt{\left[\frac{1}{D_t} \sum_{d=1}^{D_t} (u_{i,td} - ua_{it})^2 \right]} \quad (\text{A4})$$

where $u_{i,td}$ is the OLS residual from (A3), D_t is the number of trading days in quarter t and ua_{it} is the sample mean of daily idiosyncratic returns in quarter t . Thus, σ_{it} is a time-varying measure of firm-specific uncertainty that is purged from predictable market-based variation.

Finally, an aggregate (including sub-sector specific) measure(s) of uncertainty can be calculated by assuming that the firm specific uncertainty follows an autoregressive process of the following form:


$$\log \sigma_{it} = \gamma_i + \rho_i \log \sigma_{it-1} + v_t + \epsilon_{it} \quad (\text{A5})$$

where γ_i is a firm fixed effect to account for cross-sectional heterogeneity, and the aggregate level idiosyncratic volatility common to all firms is captured by the time fixed effects estimate v_t .

To calculate the uncertainty measures, a sample of 490 listed UK firms with an existing stock price quote in the Capital IQ database at the end of the sample were used. While this sample is obviously much smaller than the one used for the main analysis, this cannot be avoided, as not all the firms in the main analysis are quoted firms. The 490 firms provide a relatively large sample though, which allows for aggregation into the subsector as well as the total aggregate level.

The resulting series are shown in Chart A1, together with the aggregate RCC measure. The sectoral uncertainty measures are highly correlated, which is unsurprising, but also exhibit intuitive diversions, like the large uncertainty related to the IT sector in the early 2000s and the spike in the financial sector measure during the financial crisis. The estimates of ρ_i in (A5) range between 0.3-0.5, pointing to fairly persistent firm-level uncertainty, and also in line with the results of Gilchrist et al. (2014) for the US economy.

Chart A1: RCC and aggregated firm-level uncertainty measures



The chart shows 4-quarter moving averages of the uncertainty measure.

Appendix 2: Factor model with missing observations²⁰

The firm-level data needs to be filtered into macro-level factors; i.e., all firm-level observations for a particular variable (e.g. capital expenditure) needs to be filtered into one indicator time series, which proxies for the variable in the macro-level analysis. This is achieved with a method introduced by Giannone et al. (2008), which allows for a large dataset potentially with missing observations, to be filtered with Kalman smoothing techniques into one (or more) factors.

In the current case, the starting point for the filtering is a model for the firm-level time series of the following form:

$$x_{i,t} = \mu_i + \lambda_i f_t + \epsilon_{it} \quad (\text{A6})$$

where $t=1, \dots, T$ is time, $i=1, \dots, n$ indicates the sample of firms, μ_i is a firm-specific constant, x_{it} is the variable in question (e.g. firm-specific capital expenditure), $\lambda_i f_t$ is the factor f_t with the loading coefficient λ_i and ϵ_{it} is an error term, which are orthogonal white noise across the firms. Hence, according to (A6), x_{it} depends on a firm-specific part and a (latent) common factor across all firms.

Equation (A6) is a special case of a more general model where there can be more than one factor. For the model used in this paper, the one-factor model suffices, given that the objective is to find the one underlying factor driving the economy-wide dynamics. Furthermore, an assumption made is that the factor accounts for most of the co-movements in the variable across the firms, while the error term accounts for firm-specific idiosyncratic shocks. This exploits a basic assumption in the modelling framework; due to collinearity of the firm-specific series, a projection onto a common factor is able to capture the bulk of the dynamic interaction among the series in a parsimonious model.

Matrix form of equation (A4) is the following:

$$x_t = \mu + \Lambda F_t + \epsilon_t \quad (\text{A7})$$
$$\epsilon_t \sim \mathbb{N}(0, V_e)$$

²⁰ This appendix follows Giannone et al. (2008), which also has more details on the method. The original paper uses the method to nowcast real-time quarterly GDP from a monthly dataset with missing observations and ragged edges; the current analysis only exploits the relevant part of the original modelling framework.

where x_t is an $n \times 1$ vector of firm-specific observations, Λ is an $n \times 1$ vector of factor loadings, F_t is the common factor (which we are ultimately interested in for the analysis) and ϵ_t is multivariate white noise with a diagonal covariance matrix V_ϵ . To be able to use Kalman filtering techniques to estimate the model, the dynamics of the factor are parameterised in the following form:

$$F_t = AF_{t-1} + Bu_t \tag{A8}$$

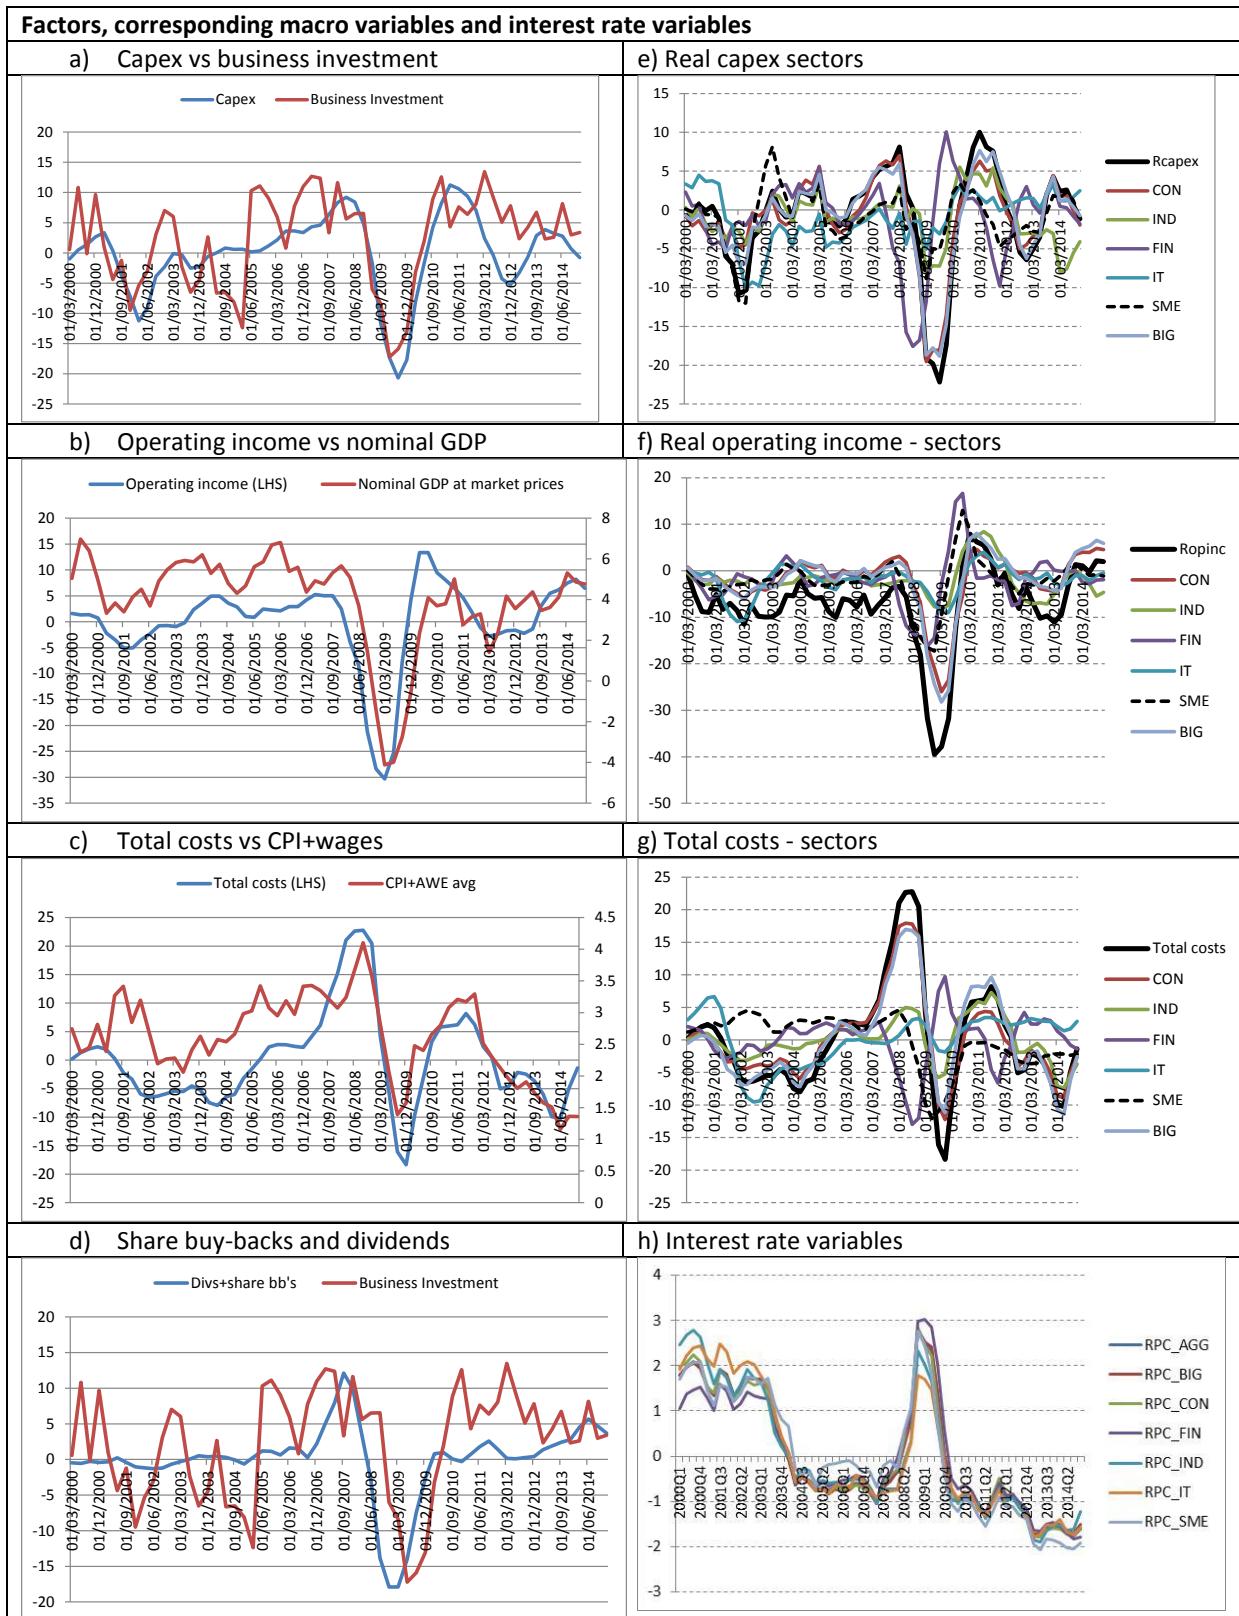
$$u_t \sim \mathcal{N}(0,1)$$

where A and B are scalars, indicating that the model has one common factor and one common shock, and u_t is a white noise shock of the factor (with a normalised variance of 1). The shock is also assumed to be orthogonal to ϵ_t .

Recent literature has shown that equations (A7)-(A8) can be estimated with a principal component analysis. However, this is not possible when there are missing observations in the dataset. Equations (A7)-(A8) represent the state space form of the model, which can be estimated with Kalman smoother, as suggested by the two-step estimator studied by Doz et al. (2006). In this estimator, the variance of the error terms (i.e., the diagonal elements of the V_ϵ matrix) is parameterised to receive its actual value when the observation is available and a value of infinity when it is missing. The Kalman smoother is then used to calculate the accuracy of both the expected value of F_t and its variance, and the algorithm places a zero weight on the missing observations, as their variances are large.

As an additional advantage of the methodology, Doz et al. (2006) have shown that the two-step estimator of the factor is consistent when both n and T are large and hence, it provides efficiency gains over principal components methods. The reason for this is that due to the law of large numbers, the idiosyncratic component becomes negligible as n increases, and hence, as long as it is confined to the idiosyncratic part of the model, any misspecification (like heteroscedasticity of error terms ϵ_{it} across firms) of the model does not compromise consistency.

Appendix 3: Additional data charts



Appendix 4²¹

This section describes the basic principles of the two sign restriction methods used in the analysis, the pure sign restriction as well as the penalty function approach.

Pure sign-restrictions approach

Let ε_t denote the ($K \times 1$) vector of structural VAR model innovations derived from equation (A1). To construct structural impulse responses, one needs an estimate of the $K \times K$ matrix C in $u_t = C\varepsilon_t$. Let $\Sigma_u = P\Lambda P$ and $C = P\Lambda^{1/2}$ such that C satisfies $\Sigma_u = CC'$. Then $C = BD$ (where B is a matrix of structural parameters obtained through a Choleski decomposition of the reduced form parameters) also satisfies $\Sigma_u = CC'$ for any orthonormal $K \times K$ matrix D .

It is possible to examine a wide range of possibilities for C by repeatedly drawing at random from the set \mathbf{D} of orthonormal rotation matrices D . Following Rubio-Ramirez et al. (2010), I construct the set \mathbf{C} of admissible models by drawing from the set \mathbf{D} of rotation matrices and discarding candidate solutions for C that do not satisfy a set of a priori sign restrictions on the implied impulse response functions. The procedure follows these steps:

1. Draw a $K \times K$ matrix K of NID(0,1) random variables. Derive the QR decomposition (to produce an orthonormal matrix and an upper-triangular matrix) of K such that $K = QR$ with the diagonal of R normalised to be positive.
2. Let $D = Q$. Compute impulse responses using the orthogonalisation $C = BD$. If all implied impulse response functions satisfy the sign restrictions, keep D . Otherwise, discard D .
3. Repeat the first two steps a large number of times, recording each D (and the corresponding impulse response functions) that satisfy the restrictions. The resulting \mathbf{C} comprises the set of admissible structural VAR models.

Penalty function approach

Define an impulse vector, which is a vector $a \in \mathbb{R}^N$ such that there exists some matrix A , where a is a column of A , such that $AA' = \Sigma_u$. Thus, the j :th column of A represents the immediate impact, or impulse vector, of a one standard error innovation to the j :th fundamental innovation, which is the j :th element of the structural error term ε_t . Furthermore, let C be the lower-triangular Choleski factor²² of

²¹ This section draws on Mountford and Uhlig (2009) and Rubio-Ramirez et al. (2010).

²² The Choleski factorisation is not used for identification here. It only serves as a computational tool, and any other factorisation would deliver the same results.



Σ_u and $Q = [q^{(1)}, \dots, q^{(s)}]$ be an $N \times S$ matrix of orthonormal rows $q^{(i)}$, where S is the number of shocks to be identified in the model. Any impulse vector can then be written as, $a = Cq$, where q is the relevant column of Q , and $q = [q_1, \dots, q_N]$, $\|q\| = 1$. Hence, q are the identifying weights to be determined. Following Uhlig (2005), the impulse responses for the impulse vector a can be written as a linear combination of the impulse responses to the Choleski decomposition of Σ_u as follows. Let $r_a(k)$ be the N -dimensional impulse response at horizon k to the impulse vector a . The linear combination can then be written as:

$$r_a(k) = \sum_{i=1}^N q_i r_i(k) \quad (\text{A7})$$

where q_i is the i :th entry of q .

Next, define the penalty function f on the real line as $f(x) = 1000x$ if $x > 0$ and $f(x) = x$ if $x \leq 0$. Let s_j be the standard error of variable j . Let $J_{s,+}$ be the index set of variables, for which identification of a given shock restricts the impulse response to be positive, and let $J_{s,-}$ be the index set of variables, for which identification restricts the impulse responses to be negative. To impose these sign restrictions, one solves for the weights q and thus $a = Cq$ by solving the following minimisation problem:

$$q = \arg \min T(Cq) \quad (\text{A8})$$

where the criterion function $T(a)$ is given by:

$$T(a) = \sum_{j \in J_{s,+}} \sum_{k=0}^K f\left(-\frac{r_{ja}(k)}{s_j}\right) + \sum_{j \in J_{s,-}} \sum_{k=0}^K f\left(\frac{r_{ja}(k)}{s_j}\right) \quad (\text{A9})$$

The criterion function thus sums the penalties over the periods $k = 0, \dots, K$ (in my case, $K=4$) following the shock and over the indices of variables with positive ($J_{s,+}$) and negative ($J_{s,-}$) sign restrictions, respectively. The impulse responses are normalised by the standard error s_j of variable j . The penalty function is, of course, arbitrary, but widely used in the literature as well as robust to changes.

To identify more than one impulse vectors $a^{(i)}$, the first vector can be identified as detailed above, after which the vector pertaining to the i :th shock can be additionally imposing orthogonality to the first shock. In my case, there are three shocks (demand, supply and cost of capital shocks), which are ordered to be causally subsequent to the pure investment (or “residual”) shock.

Appendix 5

Impulse responses with Choleski ordering (investment is ordered last)

