

Sectoral Heterogeneity in the Transmission of Monetary Policy: Panel Analysis of UK Business Investment

Veena Sasikaran

Professional Economist BSc and Apprenticeship Level 6

School of Economics

University of Kent, 2026

Executive Summary

This dissertation investigates sectoral heterogeneity in the transmission of monetary policy to business investment in the United Kingdom. Using a quarterly panel of 20 UK industries from 1997Q1 to 2025Q4, the analysis combines industry-level business investment data from the Office for National Statistics (ONS) with high-frequency monetary policy shocks from the UK Monetary Policy Event-Study Database (UKMPD). Panel local projections are estimated using the UKMPD Target factor, with sector-specific interactions testing deviations from the panel-average response and standard errors two-way clustered by industry and quarter. The aggregate response to a contractionary monetary policy shock is negative and hump-shaped, reaching -2.40% at a 15-quarter horizon, but is not statistically significant at conventional levels. The central finding is that sectoral estimates reveal substantial heterogeneity: 16 of the 20 industries reject homogeneity at the 5% level for at least one horizon. The strongest negative relative responses are concentrated in Solid fuels and oil refining, Construction, Mining and Quarrying, Other manufacturing, Metals and metal goods, and Real estate. These findings suggest that monetary policy transmission is uneven across sectors, with patterns consistent with greater sensitivity in capital-intensive, cyclical and finance-sensitive industries.

AI Statement

I acknowledge the use of generative AI in the literature search and code development in this paper. However, the work reported remains my own.

Acknowledgements

I am deeply grateful to my colleagues at the Bank of England, fellow degree apprentices, and my family for their support throughout the four years of my degree apprenticeship.

1. Introduction

A central question in macroeconomics is how monetary policy transmits to the real economy. The Bank of England (BoE) defines monetary policy as the action it takes to influence the quantity of money in the economy and the cost of borrowing, with the primary objective of maintaining low and stable inflation at a 2% target over the medium-term. Since gaining operational independence in 1997, the BoE's Monetary Policy Committee (MPC) has pursued this objective primarily through Bank Rate.

Business investment is a key channel through which the transmission of monetary policy operates, accounting for approximately 10% of UK GDP and, as one of the most cyclical components of aggregate demand, contributing disproportionately to the volatility of output (Bunn et al., 2021). Investment also matters beyond short-run demand; by increasing the capital stock, it supports productivity growth and shapes the balance between demand and supply that ultimately determines inflation. Understanding the factors driving business investment is therefore central to the conduct of monetary policy.

Monetary policy shocks, however, are unlikely to transmit uniformly to investment across the economy. Industries differ in their capital intensity, dependence on external finance, and exposure to interest-sensitive demand, all of which may shape the magnitude and timing of sectoral investment responses. Standard transmission channels provide a rationale for this heterogeneity: tighter policy raises the user cost of capital, increases external financing costs, and can weaken firms' internal funds through higher debt-servicing costs (Hall and Jorgenson, 1967; Bernanke and Gertler, 1995; Fazzari, Hubbard and Petersen, 1988).

In this context, recent empirical evidence supports the relevance of heterogeneity; Cloyne, Ferreira, Froemel and Surico (2023) show that firm-level investment responses are stronger among younger, financially constrained firms, while Choi, Willems and Yoo (2024) find larger industry-level output responses in sectors with lower collateral value and durable-goods production. While these studies suggest that monetary transmission can vary across firms and sectors, there remains significant scope for UK-specific evidence regarding industry-level business investment.

This dissertation addresses this gap by examining the dynamic response of UK industry-level business investment to monetary policy shocks. Answering this question is critical because while the BoE sets a single Bank Rate, its effects may be unevenly distributed across sectors. If certain industries bear a disproportionate share of the contractionary adjustment, aggregate impulse responses may obscure the sectoral consequences of monetary policy. Examining whether such heterogeneity exists is therefore an important first step before investigating its underlying causes.

1.1 Research Question

This dissertation therefore addresses the following question: do UK monetary policy shocks transmit heterogeneously to business investment across sectors? To answer this, the study estimates panel local projections (Jordà, 2005) on ONS sectoral business investment data covering twenty industries, using the high-frequency-identified “Target” monetary shocks of Braun, Miranda-Agrippino and Saha (2024), over the sample 1997Q1 to 2025Q4. The dissertation hypothesises that a contractionary monetary policy shock exerts a statistically significant negative effect on business investment across UK sectors, with heterogeneity in the magnitude and timing of the response across industries.

1.2 Disposition

This dissertation begins with the theoretical framework, followed by a literature review of monetary transmission to investment. The data section describes the business investment series, the monetary policy shocks, and the macroeconomic controls, alongside robustness checks to ensure the reliability of the underlying data. The methodology section then sets out the empirical approach, after which the results are presented and discussed. Finally, the dissertation situates these findings within the existing literature, acknowledges limitations, and concludes with policy implications and suggestions for future research.


2. Theoretical Framework

This section sets out the theoretical and empirical framework underpinning the analysis. It first describes the primary mechanisms through which changes in Bank Rate affect business investment, before identifying the sectoral characteristics that may generate heterogeneous responses across industries.

2.1 Monetary Transmission Framework

As outlined by BoE (2024), monetary policy transmits through six channels: inflation expectations, investment, wealth, interest rates, credit, and the exchange rate, which feed through to the real economy via aggregate demand, wages, and consumer prices. Figure 1 illustrates this framework; this dissertation focuses on the investment channel (shaded).

Figure 1



Source: adapted from Burr and Willems, 2024.

The classical exposition of this channel is provided by Tobin (1969) who establishes a general equilibrium framework in which monetary policy reaches the real economy primarily

by changing the market valuation of physical assets relative to their replacement cost. Tobin formalises this through the variable q , defined as the ratio of the market price of existing capital to the cost of producing new capital. When $q > 1$, the market values installed capital more highly than the cost of producing new capital, which stimulates investment. Conversely, when $q < 1$, as it is cheaper to acquire existing capital through the market than to produce new assets, investment is disincentivised. Monetary policy influences this mechanism by altering asset prices and returns, which in turn move q and affect investment and aggregate demand. However, the sensitivity of this adjustment depends significantly on the underlying characteristics of firms and industries.

2.2 Sources of Sectoral Heterogeneity

Three mechanisms provide a theoretical basis for expecting sectoral heterogeneity in the response of business investment to a monetary policy shock:

- i) The cost-of-capital channel implies that capital-intensive sectors are more sensitive to interest-rate changes because policy shocks affect the user cost of capital (Hall and Jorgenson, 1967).
- ii) The credit channel predicts stronger responses in sectors that rely more heavily on bank finance or contain a larger share of small firms, since monetary tightening raises the external finance premium faced by financially constrained firms (Bernanke, Gertler, and Gilchrist, 1999).
- iii) The cash flow mechanism suggests that highly leveraged sectors may reduce investment more sharply when higher debt-servicing costs weaken internal funds (Fazzari, Hubbard, and Petersen, 1988).

Together, these mechanisms imply that monetary policy shocks may generate heterogeneous investment responses across industries. This dissertation therefore estimates aggregate and sectoral impulse responses to UK monetary shocks using panel local projections and interprets the observed heterogeneity considering the channels outlined above.

3. Literature Review

This section critically reviews the theoretical and empirical literature on the transmission of monetary policy to investment and focuses on heterogeneity across industries and firms,

providing a basis for the empirical analysis conducted in this dissertation. It concludes by identifying an unresolved tension between firm and industry-level research on transmission heterogeneity, which motivates the empirical contribution of this paper.

3.1 Theoretical Channels of Monetary Policy Transmission

A large theoretical literature has studied the channels through which monetary policy affects real activity, of which the cost-of-capital, balance-sheet, and credit channels are most directly relevant for business investment. While these channels share the prediction that a monetary tightening reduces investment, they differ in their cross-sectional pattern mechanisms. Each predicts that a different industry characteristic: capital intensity, balance-sheet strength, or credit dependence should determine the size of the response. This distinction is central to the empirical strategy adopted in this dissertation.

3.1.1 Cost of Capital Channel

The cost-of-capital, or interest rate, channel originates in the neoclassical theory of investment developed by Jorgenson (1963), in which firms determine their desired capital stock by equating the marginal product of capital to its user cost. A monetary tightening raises real interest rates, increasing the user cost and reducing the investment needed to reach the target capital stock.

A complementary mechanism, drawn from the microeconomic literature on investment under uncertainty, operates through the irreversibility of capital. Under partial irreversibility, Dixit and Pindyck (1994) show that firms hold a real option to delay whose value is increasing in the volatility of expected returns, raising the threshold return required to undertake investment. Industries with durable, specialised capital therefore face two reinforcing mechanisms. Higher real rates reduce the desired capital stock through the user-cost channel, while irreversibility raises the option value of waiting and increases the hurdle rate for new investment projects. Both imply that capital-intensive industries should exhibit the largest and most persistent contractions in investment following a monetary tightening.

3.1.2 Balance Sheet Channel

Bernanke and Gertler (1989, 1995) show that, in the presence of agency frictions between borrowers and lenders, firms face an external finance premium - the additional cost of

external funds relative to internal finance. This premium depends negatively on borrower net worth, liquidity, cash flow and collateral. A monetary tightening weakens firm balance sheets through three key mechanisms:

- i) it raises debt-servicing costs;
- ii) reduces cash flow;
- iii) lowers asset values used as collateral.

As creditworthiness deteriorates, external finance premiums rise and the effective cost of borrowing increases, amplifying the contraction in investment beyond what the user-cost channel alone implies. Bernanke, Gertler and Gilchrist (1999) formalise this dynamic as the financial accelerator. Empirical support for the channel comes from Gertler and Gilchrist (1994), who find that smaller US manufacturing firms, proxying for financial constraints, contract more sharply than larger firms following a monetary tightening.

3.1.3 Credit Channel

The credit, or bank lending, channel emphasises the role of financial intermediaries in propagating monetary policy. Bernanke and Blinder (1988, 1992) develop the canonical model, in which a monetary tightening reduces the supply of bank loans, with the strongest effects on bank-dependent borrowers who lack access to alternative sources of finance such as bond or equity markets.

Empirical support comes from Kashyap and Stein (2000), who show using US bank-level data that the lending response is concentrated among smaller and less-liquid banks. Like the balance-sheet channel, the bank lending channel is rooted in asymmetric information between borrowers and lenders but operates on a different margin: the balance-sheet channel works through the demand side, via borrower net worth, whilst the bank lending channel works through the supply side, via the willingness of banks to extend credit. Bernanke and Gertler (1995) characterise the two together as a "broad credit channel" amplifying the conventional interest-rate channel.

3.2 Empirical Evidence on Aggregate Investment Responses

Identified monetary policy shocks are generally found to aggregate investment, with contraction building over time before gradually unwinding. Christiano, Eichenbaum and Evans (1999, 2005) use recursively identified VAR evidence for the US and find delayed,

hump-shaped real responses to monetary policy shocks. In their framework, output, consumption and investment respond gradually rather than immediately, with peak effects occurring after roughly one to two years. Peersman and Smets (2003) document similar timing and shape for the euro area, while Ramey (2016) shows that this broad profile is robust across countries and identification methods, although the magnitude varies. However, recursive VAR identification relies on contemporaneous timing restrictions and may recover shocks that contain central bank information effects (Miranda-Agrippino and Ricco, 2021). This motivates the use of high-frequency monetary policy shocks in this dissertation.

The evidence addressing the UK remains comparatively limited. This scarcity derives largely from the difficulty of identifying exogenous policy changes in a small open economy impacted by global exchange rate and trade shocks. Cloyne and Hürtgen (2016) construct a narrative UK monetary policy shock series, following Romer and Romer (2004), and find that a contractionary shock reduces UK output, with the largest effect occurring after roughly two years. Cloyne et al. (2023) use a panel LP-IV approach and show that firm-level investment also falls significantly after monetary tightening, peaking between the second and third year. Together, these studies suggest UK monetary transmission is broadly consistent with results found in the aggregate US and euro-area literature.

3.3 Identification of Monetary Policy Shocks

To assess the causal effects of monetary policy, it is necessary to isolate exogenous policy changes from interest-rate movements that reflect the central bank's response to economic conditions. A monetary policy shock is therefore the unexpected component of a policy change, making identification central to the monetary transmission literature.

Early studies identified shocks using recursively identified VARs, building on Sims (1980) and applied to monetary transmission by Bernanke and Blinder (1992) and Christiano, Eichenbaum and Evans (1999). While influential, this approach relies on restrictive timing assumptions, particularly the contemporaneous ordering imposed by the Cholesky decomposition. Small VARs may also omit information available to policymakers, raising concerns that policy-rate residuals are not true structural shocks (Sims, 1992; Miranda-Agrippino and Ricco, 2021).

Subsequent work addressed these concerns by using richer information sets and narrower identification windows. Narrative approaches, such as Romer and Romer (2004), construct monetary policy shocks by removing the Federal Reserve's response to its own forecasts, while Cloyne and Hürtgen (2016) adapt this approach to the UK using Bank of England forecasts. These methods improve on small VARs by conditioning directly on policymakers' information, though they are typically available at a lower frequency. Alternatively, high-frequency identification exploits movements in financial market prices within narrow windows around policy announcements, treating these changes as monetary policy surprises (Kuttner, 2001; Gürkaynak, Sack and Swanson, 2005; Gertler and Karadi, 2015). This approach offers sharper timing and is better suited to tracing the dynamic effects of policy shocks.

Recent work extends the high-frequency approach to the UK. Cesa-Bianchi, Thwaites and Viccondoa (2020) provide an early UK high-frequency shock series using surprises in short sterling futures. Braun, Miranda-Agrippino and Saha (2024) contribute to the UK identification literature by constructing the UK Monetary Policy Event-Study Database, which measures intraday surprises around MPC announcements and Monetary Policy Report press conferences from 1997 to 2023. Building on Gürkaynak, Sack and Swanson (2005) and Swanson (2021), they decompose these surprises into Target, Path, and QE factors, while using the Jarociński and Karadi (2020) sign restriction to separate monetary policy shocks from central bank information effects.

3.4 Sectoral Heterogeneity in Monetary Policy Transmission

A growing literature shows that aggregate responses to monetary policy conceal substantial heterogeneity across firms and industries. Ganley and Salmon (1997) provide the main UK benchmark, estimating recursively identified VARs of output for twenty-four industries over 1970-1995. They document sizeable cross-industry dispersion following a monetary tightening, with construction and durable manufacturing responding more strongly than agriculture and food-related sectors. Comparable evidence is provided by Dedola and Lippi (2005) and Peersman and Smets (2005), who find larger effects in capital-intensive and durable-goods industries.

More recently, Choi, Willems and Yoo (2024) extend the industry-level literature using a panel of 102 countries and 22 industries. Interacting monetary policy shocks with industry characteristics, they find stronger output effects in industries with less-collateralisable assets and smaller firms, consistent with the credit channel, and in durable-goods industries, consistent with the interest-rate channel. They find no support for the cost channel and only limited support for an export channel.

This literature establishes that monetary transmission differs across industries, but important limitations remain. Much of the earlier evidence relies on recursive VAR identification - which, as discussed, may fail to isolate truly exogenous shocks and relatively short historical samples. In addition, industry characteristics such as capital intensity, durability, firm size, and financial structure are highly correlated, complicating efforts to isolate individual competing transmission mechanisms. While recent cross-country work provides broader evidence on channels, it focuses primarily on output and prices rather than the granular response of business investment. This leaves significant scope for updated UK evidence using modern identification strategies to assess whether industry-level investment responds heterogeneously to monetary policy shocks.

3.5 Contribution and Relation to the Literature

This dissertation contributes to the literature by updating the UK industry-level evidence on monetary policy transmission, with a specific focus on business investment. While recent firm-level studies demonstrate that monetary policy affects firms differently, and industry-level research identifies cross-country variation in output responses, less is known about whether such heterogeneous responses are detectable in UK sectoral investment data. The analysis utilises the Target shock series from Braun, Miranda-Agrippino, and Saha (2024), panel local projections, and UK industry investment data. It documents the distribution of investment responses across industries and tests whether individual sectors differ statistically from the panel-average response. By bridging older UK industry-level evidence with recent firm-level and cross-country research, this dissertation assesses whether monetary policy heterogeneity is detectable at the sectoral level, a level of aggregation highly relevant for policy analysis.

4. Data

This dissertation combines sector-level business investment data from the Office for National Statistics (ONS) with high-frequency monetary policy shocks from the UK Monetary Policy Event-Study Database (UKMPESD; Braun, Miranda-Agrippino and Saha, 2024). The resulting quarterly industry panel covers 20 UK industries over 1997Q1-2025Q4, totalling 2,320 observations. The sample begins in 1997Q1 to align the analysis with the post-independence monetary policy regime established by the Bank of England Act 1998, under which the MPC was given operational responsibility for setting Bank Rate to achieve price stability. This restricted sample provides a more consistent monetary policy framework and avoids pooling observations from different regimes.

4.1 Investment Data

The objective of this dissertation is to assess whether monetary policy shocks have heterogeneous effects across UK industries. The primary dependent variable is real business investment by industry, obtained from the ONS Business Investment by Industry and Asset dataset. The series are measured as chained volume measures and are seasonally adjusted. Industry classifications follow SIC 2007, identified by ONS CDIDs. The final dataset covers 20 industries across manufacturing, non-manufacturing production, construction, distribution, and services. Figure 2 plots aggregate UK business investment, constructed as the sum of the 20 industry-level series:

Figure 2




Notes: Aggregate UK business investment is constructed as the sum of 20 industry-level chained-volume, seasonally adjusted series, 1997Q1–2025Q4. The shaded area denotes the COVID-19 indicator period, 2020Q1–2022Q3.

Two features of the data inform the empirical specification. First, investment exhibits a strong upward trend in levels, rising from £37.3bn in 1997Q1 to £75.2bn in 2025Q4, while also displaying substantial cyclical variation. This motivates transforming the series before estimation. Second, the COVID-19 period is associated with an exceptional collapse and rebound in investment, with aggregate investment falling by around 22% between 2019Q2 and 2020Q2. Since this volatility reflects pandemic-related disruption rather than monetary policy transmission, the baseline specification includes a COVID-19 indicator. This reduces the risk that estimated responses to monetary policy shocks are confounded by pandemic-specific disruption.

4.2 Monetary Policy Shock Data

Monetary policy shocks are sourced from the UKMPESD, which constructs high-frequency measures of unexpected policy changes using financial market data around Bank of England Monetary Policy Committee (MPC) announcements. The analysis focuses on the Target shock, which captures surprises to near-term policy rates and loads most heavily on the short end of the yield curve. Positive values are interpreted as contractionary monetary policy surprises. Figure 3 plots the quarterly aggregated Target shock over the sample period:

Figure 3



Notes: Quarterly aggregates of MPC-day Target factor surprises from the UKMPESD. Announcement-level surprises are summed within each calendar quarter. Positive values denote contractionary surprises to the current Bank Rate. The shock is standardised to unit standard deviation in estimation.

Announcement-level shocks are aggregated to a quarterly frequency by summing the shock values within each calendar quarter, aligning the data with the frequency of the dependent variable. The series fluctuates around zero with no evident trend, consistent with its construction as a measure of unanticipated monetary policy shocks. The identifying assumption underlying its use is that movements in the Target factor within the narrow announcement window primarily reflect monetary policy news rather than other macroeconomic information released contemporaneously.

4.3 Macroeconomic Controls

The specification includes a set of quarterly macroeconomic controls to account for aggregate conditions that may jointly affect business investment and the stance of monetary policy. Real GDP growth accounts for the business cycle and aggregate demand conditions, while CPI inflation captures price pressures relevant to monetary policy decisions and real financing conditions. The unemployment rate controls for labour market slack, which may influence firms' expectations and investment plans. All control variables are obtained from the ONS, measured at a quarterly frequency, and enter the baseline specification with lags to reduce simultaneity concerns.

4.4 Summary Statistics

Table 1 reports summary statistics for the variables used in the sample. Mean industry-quarter investment is £2.8bn, with substantial dispersion reflecting differences in sector size. The Target shock has a mean close to zero and a quarterly standard deviation of 0.074 percentage points, consistent with its construction as a monetary policy surprise.

Table 1

Variable	Mean	Std. Dev	Minimum	Maximum	Count
Investment (£m, CVM SA)	2,763.3	2,693.5	34.000	18,030.0	2,319
Quarterly investment growth (%)	1.488	15.138	-84.589	144.9	2,299
Target shock (quarterly aggregate, p.p.)	-0.004	0.074	-0.245	0.247	116
Real GDP growth, annualised q/q (%)	1.818	4.155	-24.425	22.696	116
CPI inflation, y/y (%)	2.454	1.973	0.000	10.700	116
Unemployment rate (%)	5.520	1.300	3.700	8.400	116

Notes: The panel contains a maximum of 2,320 industry-quarter observations. One negative chained-volume observation, Transportation and storage in 2012Q3, is treated as missing before taking logs, leaving 2,319 non-missing investment observations.

5. Methodology

Monetary policy affects the economy with long and variable lags. BoE estimates suggest monetary policy takes around 18-24 months to have its full effect on the economy, requiring the MPC to form judgments over horizons extending several years ahead. For business investment, empirical literature suggests this lag is longer, with peak responses materialising two to three years after a shock (Cloyne et al., 2023). Capturing this adjustment path requires an estimator that traces responses over multiple quarters rather than estimating a single average effect.

5.1 Local Projections

Accordingly, this dissertation uses local projections, following Jordà (2005), to estimate the response of UK sectoral business investment to identified monetary policy shocks. Local projections trace the dynamic response of an outcome variable to a shock at each horizon following the shock, generating an impulse response function (IRF) that captures both the magnitude and the timing of the effect. The approach is well suited to this dissertation because it allows impulse responses to be estimated separately at the aggregate level and for each industry, providing a framework for comparing sectoral responses against an aggregate benchmark. They avoid the cross-equation restrictions of a VAR, so impulse responses are robust to misspecification of the underlying dynamic system (Plagborg-Møller and Wolf, 2021). They also accommodate panel data through industry fixed effects, in contrast to the

sector-by-sector VARs of Ganley and Salmon (1997) and extend naturally to the interaction specifications used in the heterogeneity tests of Section 4.4.

For each horizon $h = 0, 1, \dots, H$, local projections estimate a separate OLS regression of cumulative investment growth between $t - 1$ and $t + h$ on the monetary policy shock at time t , together with fixed effects and relevant lags. The coefficient on the shock at each horizon provides the impulse response at that horizon, so the sequence of OLS coefficients traces the path of investment following a monetary policy shock.

5.2 Baseline Panel Specification

The baseline specification estimates the average sectoral investment response to a monetary policy shock. For each horizon $h = 0, 1, \dots, 20$, the following regression is estimated on the panel of $i = 1, \dots, 20$ industries, observed over $t \in [1997Q1, 2025Q4]$:

$$y_{i,t+h} - y_{i,t-1} = \alpha_i + \beta_h z_t + \Gamma'_h W_{i,t-1} + \lambda_h D_t^{COVID} + \varepsilon_{i,t+h} \quad (1)$$

where $y_{i,t} = \ln(Investment_{i,t})$ for industry i at quarter t . The dependent variable $y_{i,t+h} - y_{i,t-1}$ represents the cumulative log-change in investment between the quarter immediately preceding the shock and horizon h ($t - 1$ to $t + h$). The coefficient of interest, β_h , is the percent response of investment h quarters after a one standard-deviation tightening surprise. Coefficients are interpreted as percentage responses ($100 \times \beta_h$). The standardised quarterly Target shock is represented by z_t .

The control vector $W_{i,t-1}$ collects p lags of the monetary policy shock, lagged investment growth, and lagged macroeconomic controls:

$$W_{i,t-1} = \begin{pmatrix} z_{t-1} \\ \vdots \\ z_{t-p} \\ \Delta y_{i,t-1} \\ \vdots \\ \Delta y_{i,t-p} \\ \Delta \ln \text{GDP}_{t-1} \\ \vdots \\ \Delta \ln \text{GDP}_{t-p} \\ \pi_{t-1} \\ \vdots \\ \pi_{t-p} \\ u_{t-1} \\ \vdots \\ u_{t-p} \end{pmatrix}$$

where $\Delta y_{i,t-k}$ represents log investment growth at lag k , $\Delta \ln \text{GDP}_t$ represents year-on-year real GDP growth, π_t represents CPI inflation, and u_t represents the unemployment rate. Macroeconomic controls enter at lags $k \geq 1$, avoiding conditioning on contemporaneous variables that may respond to z_t .

Industry fixed effects are represented by α_i absorb time-invariant heterogeneity in average investment growth across sectors. Additionally, a COVID-19 indicator, D_t^{COVID} , is included to capture the exceptional economic disruptions, shown in Figure 1, which are unlikely to operate through standard monetary transmission mechanisms; this indicator is equal to one for the period 2020Q1-2022Q3. The error term is represented by $\varepsilon_{i,t+h}$.

5.3 Hypotheses

The empirical analysis tests whether monetary policy shocks have homogeneous effects across UK industries. The hypotheses are:

- H_0 : Monetary policy shocks have homogenous effects on business investment across UK industries
- H_1 : Monetary policy shocks have heterogenous effects on business investment across UK industries

Formally, this corresponds to testing whether the industry-specific differential response is equal to zero:

$$H_0: \theta_{h,j} = 0 \quad (2)$$

against

$$H_1: \theta_{h,j} \neq 0 \quad (3)$$

Where $\theta_{h,j}$ denotes the differential response of industry j at horizon h . Rejecting H_0 implies that the investment response of industry j differs statistically from the benchmark sectoral response at that horizon.

5.4 Sectoral Heterogeneity Specification

To test these hypotheses, the baseline local projection is augmented with a sector-specific interaction between the monetary policy shock and an industry indicator. For each industry $j = 1, \dots, 20$ and each horizon h , the following equation is estimated:

$$y_{i,t+h} - y_{i,t-1} = \alpha_i + \beta_h z_t + \theta_{h,j}(z_t \cdot \mathbb{1}[i = j]) + \Gamma'_h W_{i,t-1} + \lambda_h D_t^{COVID} + \varepsilon_{i,t+h} \quad (4)$$

where $\mathbb{1}[i = j]$ is an indicator equal to one when observation i belongs to industry j . The average response across the remaining nineteen industries is represented by β_h and the differential response of industry j relative to that benchmark is represented by $\theta_{h,j}$. The total response of industry j at horizon h is $\beta_h + \theta_{h,j}$. Equation (4) is estimated separately for each industry, with standard errors two-way clustered by industry and quarter. The hypothesis is tested coefficient-by-coefficient whereby rejection of the null at any sector-horizon pair (j, h) constitutes evidence against homogeneous transmission of monetary policy across UK industries, with emphasis placed on systematic patterns across horizons rather than isolated significant coefficients.

5.5 Identification, Estimation, and Inference

Identification exploits exogenous variation in monetary policy captured by high-frequency surprises around MPC announcements. These shocks isolate the unanticipated component of policy rate changes and are therefore treated as orthogonal to contemporaneous macroeconomic conditions; under this assumption, the estimated impulse responses support a causal interpretation.

Equation (1) and (4) are estimated by ordinary least squares (OLS) separately at each horizon. Due to the high-frequency Target factor being treated as a direct measure of the monetary policy shock (rather than as an instrument for changes in Bank Rate), OLS is preferred over IV local projections. An IV approach would be more appropriate if the shock series were used as an external instrument for an endogenous policy variable (Stock and Watson, 2018). The trade-off is that OLS gives a more direct estimate, reliant on the Target factor being a valid measure of exogenous monetary policy news.

The baseline specification includes $p = 4$ lags, corresponding to one year of quarterly data. Alternative lag lengths $p \in \{2, 8\}$ are reported as robustness checks in Section 5.2. The horizon is set to $H = 20$ quarters to capture the full adjustment path of investment to monetary policy shocks.

Standard errors are two-way clustered by industry and quarter following Cameron, Gelbach, and Miller (2011). This accommodates within-industry serial correlation in $\varepsilon_{i,t+h}$, which arises by construction at horizons $h \geq 1$ where overlapping observations induce correlated residuals, alongside contemporaneous correlation across industries within a quarter from aggregate shocks not fully absorbed by the lagged controls.

5.6 Diagnostics and Robustness

Three diagnostic tests inform the empirical specification: non-stationarity in series, autocorrelation in the local-projection errors, and heteroskedasticity across industries and over time.

5.6.1 Stationarity

Regressions involving non-stationary variables may yield spurious results (Hill et al., 2008). Sector-level Augmented Dickey-Fuller tests indicate that log investment is non-stationary in levels but stationary in first differences across all twenty industries. Im-Pesaran-Shin panel unit-root tests provide complementary evidence against a unit root in the transformed investment series. The dependent variable is therefore defined as the cumulative log change in investment relative to $t - 1$, reducing the risk of spurious regression. The remaining variables enter in forms appropriate for the specification. Real GDP growth enters as a four-quarter log change. CPI inflation and unemployment enter as percentage-point rates and are included only as lagged controls. The quarterly Target shock fluctuates around zero with no persistent trend, consistent with its construction as a monetary policy surprise.

5.6.2 Autocorrelation and Heteroskedasticity


Autocorrelation and heteroskedasticity are expected in panel local projections, given the dependent variable overlaps across horizons. The baseline and heterogeneity specifications therefore use two-way clustered standard errors by industry and quarter, which are robust to arbitrary heteroskedasticity and autocorrelation within both dimensions (Cameron and Miller, 2015; Abadie et al., 2023). Clustering by industry allows for serial correlation within sectors, while clustering by quarter allows for contemporaneous correlation across industries arising from common aggregate shocks. Per-industry local projections use HC1 heteroskedasticity-robust standard errors, as two-way clustering is not feasible within single-industry regressions. The main inference on heterogeneity is instead drawn from the panel interaction specifications with two-way clustered standard errors.

6. Results

6.1 Aggregate Investment Response

Figure 4 presents the impulse response of UK sectoral investment to a one standard-deviation Target tightening, estimated from equation (1) with industry fixed-effects, four lags of the shock, lagged macroeconomic controls, and a COVID-period indicator. Standard errors are two-way clustered by industry and quarter.

Figure 4



Notes: The solid line reports the cumulative impulse response of sectoral investment to a one-standard-deviation Target-factor tightening. Dark and light shaded areas denote 68% and 90% confidence intervals, respectively.

The response is contractionary over most horizons. On impact, $\hat{\beta}_0 = 0.07\%$, close to zero. The contraction builds gradually over the first year, reaching -1.4% at $h = 4$, and remains between approximately -1.5 and -2.4% over the medium-run window $h = 8-16$, corresponding to two to four years after the shock. The largest point estimate occurs at $h = 15$, where $\hat{\beta}_{15} = -2.4\%$ ($SE = 1.57$). The response partly unwinds thereafter, reaching -0.74% by $h = 20$. The 90 per cent confidence band does not exclude zero at any horizon, so the aggregate response is not statistically significant at conventional levels. This reflects the uncertainty associated with horizon-specific local projections and volatile sectoral investment data. Table 2 reports point estimates and clustered standard errors at key horizons.

Table 2


Lag length	h = 0	h = 4	h = 8	h = 11	h = 12	h = 13	h=15	h = 16	h = 20
p = 2	0.07 (0.31)	-0.80 (0.74)	-0.81 (1.09)	-1.17 (1.24)	-1.23 (1.39)	-1.33 (1.39)	-1.70 (1.42)	-1.33 (1.37)	-0.20 (1.68)
p = 4	-0.07 (0.41)	-1.41 (0.96)	-1.60 (1.55)	-1.98 (1.81)	-2.08 (1.55)	-2.03 (1.59)	-2.40 (1.57)	-1.85 (1.57)	-0.74 (1.67)
p = 8	0.27 (0.57)	-0.18 (1.54)	0.94 (2.09)	-1.33 (1.56)	-1.82 (1.76)	-2.28 (1.62)	-1.62 (1.62)	-1.38 (1.65)	2.22 (2.16)

Notes: Entries are point estimates in per cent, with two-way clustered standard errors in parentheses.

6.2 Robustness to Lag Length

The baseline specification fixes the lag order at $p = 4$, corresponding to one full year of quarterly macroeconomic history. As a robustness check, Figure 5 re-estimates the aggregate response using $p = 2$ and $p = 8$. The $p = 2$ and $p = 4$ specifications both imply a negative medium-run response, reaching troughs of -1.70% and -2.40% at $h = 15$, respectively. The $p = 8$ response reaches -2.28% at $h = 13$, but is less persistent and turns positive by $h = 20$. The lag-robustness exercise suggests that the contractionary medium-run profile is not driven solely by the baseline lag choice, but the timing and persistence of the aggregate response are sensitive to the lag specification.

Figure 5



Notes: Cumulative IRFs of UK sectoral business investment to a one-standard-deviation Target shock under lag orders $p = 2, 4$ and 8 . Shaded areas show 90 per cent confidence intervals.

6.3 Industry-level Investment Response


While Section 5.1 reported the average dynamic response of UK investment to a Target shock, assuming a common response across all sectors, this restriction may mask important heterogeneity in the transmission of monetary policy. To assess this heterogeneity, the local projection is estimated separately for each of the 20 industries, maintaining the same baseline lag length $p = 4$. Because these regressions are estimated separately for each industry, the two-way clustering used in the panel specification is not feasible; standard errors are therefore computed using heteroskedasticity-robust HC1 estimators.

Figure 6 reports 20 per-industry impulse responses, over horizons $h = 0$ to $h = 20$. This figure allows the dynamic profile of each sector to be inspected directly, rather than imposing the common response estimated in the aggregate panel specification. Each panel shows the cumulative response of investment in one industry to a one standard-deviation Target-factor tightening. The horizontal line denotes zero, so estimates below the line indicate a contraction in investment relative to the pre-shock period, while estimates above the line indicate a positive response.

Figure 6 shows pronounced cross-industry variation in both the magnitude and persistence of the estimated responses. Mining and Quarrying, Solid fuels and oil refining, Construction, Metals and metal goods, Other manufacturing, and Chemicals and man-made fibres display negative medium-run responses. These sectors are relatively capital-intensive or cyclically exposed, so a stronger contraction is consistent with monetary transmission through higher financing costs, cash-flow constraints, and weaker expected demand.

By contrast, Engineering and vehicles, and Textiles, clothing, leather and footwear record positive average responses over the medium-run window. These estimates should be interpreted cautiously. They are best understood as evidence that these sectors respond less negatively than the average sector, or that sector-specific investment dynamics and lumpy capital expenditure are important in shaping the estimated response, rather than as evidence that monetary tightening directly raises investment.

Figure 6




Notes: Each panel reports the IRFs of investment in each industry to a one standard-deviation Target-factor tightening. Shaded areas denote 90% HC1 confidence intervals.

The confidence intervals around the per-industry estimates are wide, particularly for Transportation and storage, Hotels and restaurants, and Textiles. The individual industry impulse responses should therefore be interpreted as descriptive evidence of heterogeneity, rather than as formal statistical tests of whether one sector differs from another. The formal heterogeneity tests in Section 5.4 provide the main evidence on whether sectoral responses differ significantly from the panel average.

Figure 7 ranks the twenty industries by their average estimated investment response over horizons $h = 8$ to 16. This window corresponds to two to four years after the shock and captures the period in which the aggregate response is most contractionary. The vertical line marks zero; points to the left indicate an average contraction in investment, while points to the right indicate a positive average response.

Figure 7



Notes: Horizontal bars show approximate 90 per cent intervals based on pooled standard errors.


Mining and Quarrying (-5.73%), Solid fuels and oil refining (-3.98%), Other manufacturing (-3.74%) and Metals and metal goods (-3.54%) anchor the bottom of the distribution — sectors that tend to be capital-intensive, energy-intensive, or cyclically sensitive, exposing investment to higher financing costs and weaker expected demand. Textiles, clothing, leather and footwear (4.10%), and Engineering and vehicles (1.20%) anchor the top, with Food, drink and tobacco (-0.14%), suggesting partial insulation from the average contraction or sector-specific investment dynamics.

6.4 Sectoral Heterogeneity Tests

To test differential transmission, equation (4) is estimated where the coefficient $\theta_{h,j}$ measures whether sector j 's response at horizon h differs from the pooled response of the remaining sectors. The specification is estimated separately for each sector-horizon pair, yielding 420 single-interaction panel regressions with two-way clustered standard errors. Negative values imply that sector j contracts more than the rest of the panel; positive values imply that it contracts less or responds more positively. Figure 8 presents $\hat{\theta}_{h,j}$ estimates as a heatmap. Red cells indicate stronger contractions than the rest of the panel, blue cells indicate weaker contractions or more positive, and white cells indicate little relative difference.

Within the main transmission window ($h = 4-16$), 16 of the 20 industries reject homogeneity at the 5% significance level for at least one horizon, and 14 reject at the 1% level. The red band at the bottom of Figure 8 shows the strongest negative relative responses, concentrated in Solid fuels and oil refining, Construction, Mining and Quarrying, Other manufacturing, Real estate, and Metals and metal goods. By contrast, the blue band at the top, anchored by Textiles, clothing, leather and footwear, and Engineering and vehicles, indicates persistently positive relative responses and weaker-than-average contractions.

Figure 8



Notes: Cells report $\hat{\theta}_{h,j}$ and the pooled response of the remaining sectors at horizon h . Red cells indicate stronger contractions than the panel; blue cells indicate weaker contractions or more positive responses. Stars denote statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

To summarise the heterogeneity results, Table 3 reports selected estimates of $\hat{\theta}_{h,j}$ at $h = 4, 8, 12, 16$. These horizons capture the main transmission window without reporting all 420 sector-horizon estimates.

Table 3:

Industry	h = 4	h = 8	h = 12	h = 16
Solid fuels & oil refining	-0.71 (0.57)	-3.65** (1.59)	-5.21*** (1.36)	-5.86*** (2.12)
Real estate, renting and business	-1.57*** (0.56)	-2.98** (1.31)	-1.71 (1.40)	-0.45 (2.17)
Electricity, gas and water	-1.93*** (0.73)	-0.56 (1.00)	-1.42 (1.06)	-1.25 (1.68)
Construction	-0.04 (0.44)	-1.26* (0.71)	-0.18 (0.68)	-1.92*** (0.60)
Information and communication	-1.31** (0.56)	-1.02 (0.99)	-0.66 (0.86)	0.23 (0.81)
Agriculture, forestry and fishing	-0.56 (0.53)	-0.95 (0.67)	0.94 (0.76)	-2.21*** (0.70)
Financial intermediation	0.09 (1.00)	-0.90 (2.07)	0.37 (2.74)	0.48 (3.68)
Mining and Quarrying	0.41 (0.72)	-0.05 (1.58)	0.16 (1.94)	-2.65 (1.93)
Metals and metal goods	0.28 (0.52)	-0.74 (0.65)	-1.36* (0.70)	0.96 (0.83)
Other manufacturing	-1.51 (1.10)	1.55 (1.59)	-1.76 (1.69)	1.18 (1.46)
Distribution Services	0.45 (0.59)	-0.30 (1.09)	-2.19 (1.56)	0.13 (1.44)
Chemicals and man made fibres	0.10 (0.70)	-0.60 (0.84)	-0.66 (1.16)	3.08*** (1.18)
Education	0.11 (0.54)	0.80 (0.71)	0.41 (0.74)	-2.42*** (0.63)
Food, drink and tobacco	0.03 (0.41)	1.16* (0.65)	-0.95 (0.79)	-1.02 (0.99)
Other services	0.20 (0.69)	0.00 (0.88)	-1.07 (1.18)	4.16*** (1.41)
Health and social work	1.20 (1.15)	0.55 (1.29)	0.83 (1.37)	-1.13 (1.54)
Hotels and restaurants	0.66 (0.65)	0.11 (0.82)	1.40 (1.30)	0.74 (1.54)
Transportation and storage	-0.29 (0.50)	0.98 (0.68)	1.80*** (0.66)	0.66 (0.60)
Engineering and vehicles	1.38*** (0.39)	0.90** (0.46)	2.74*** (0.77)	3.24*** (1.13)
Textiles, clothing, leather and footwear	3.04** (1.23)	6.96*** (1.05)	8.51*** (1.29)	4.03** (1.81)

Notes: Sectors are ordered by average $\hat{\theta}_{h,j}$ over $h = 4$ to 16, most negative first. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The strongest negative relative response is for Solid fuels and oil refining, where the contraction deepens over time from -3.65 at $h = 8$ to -5.21 at $h = 12$ and -5.86 at $h = 16$, all

significant at conventional levels. This deepening profile is consistent with the long capital-renewal cycles characteristic of the sector. Construction shows a faster response: -1.57 at $h = 4$ and -2.98 at $h = 8$, both significant, before unwinding. Mining and Quarrying displays a similar early response, -1.93 at $h = 4$, significant at the 1% level, and Other manufacturing contracts significantly more than the panel average at $h = 16$ (-1.92). Conversely, Textiles, and Engineering and vehicles respond significantly less negatively than the rest of the panel: Textiles reaches 8.51 percentage points at $h = 12$, and Engineering and vehicles reaches 3.24 at $h = 16$.

6.5 Robustness, Time Fixed Effects

The baseline heterogeneity specification in Section 5.4 controls for aggregate conditions through lagged macro variables. To verify that the cross-industry pattern documented in Section 5.4 is not driven by the choice of macro controls, the same per-sector heterogeneity tests are re-estimated using two-way (industry and quarter) fixed effects in place of the macro controls. The main heterogeneity result is robust. The Pearson correlation between average sectoral heterogeneity coefficients in the baseline and time fixed effects specifications is 0.89, and the sign of the average relative response is unchanged for 85% of sectors. This indicates that the evidence of differential sectoral transmission is not driven solely by the treatment of common quarter-level shocks.

7. Conclusion

This dissertation examined whether UK business investment responds heterogeneously across sectors to monetary policy shocks.

7.1 Key Findings and Discussion

The central finding is that UK business investment responds unevenly to monetary policy shocks across sectors. The aggregate response to a Target-factor tightening follows a contractionary, hump-shaped profile, reaching -2.40% at $h = 15$, but remains statistically insignificant. This suggests that the common-response model is too restrictive given the average effect is estimated by pooling industries with different structural exposure to interest rates, credit conditions and demand cycles. This profile is consistent with the wider literature showing delayed investment responses to monetary tightening, although the peak occurs slightly later than in firm-level evidence from Cloyne et al. (2023).

The sectoral evidence provides clearer support for heterogeneous transmission. Over the main transmission window, 16 of the 20 industries reject homogeneity at the 5% level for at least one horizon. The strongest negative relative responses are concentrated in Solid fuels and oil refining, Construction, Mining and Quarrying, Other manufacturing, Metals and metal goods, and Real estate. These sectors are typically capital-intensive, cyclical, finance-sensitive, or reliant on long-lived investment projects, making them more exposed to higher user costs of capital, tighter credit conditions, weaker cash flow and lower expected demand following contractionary monetary policy. This is consistent with Ganley and Salmon's UK evidence that monetary policy shocks have uneven effects across industry output, and with Choi, Willems and Yoo's cross-country sectoral evidence. This dissertation extends that literature by showing that similar heterogeneity is also visible in UK business investment, using high-frequency identified monetary policy shocks.

This divergence is consistent with standard monetary transmission mechanisms. The strongest negative relative responses are concentrated in Solid fuels and oil refining, Construction, Mining and Quarrying, Other manufacturing, Metals and metal goods, Chemicals, and Real estate. These industries are typically capital-intensive, cyclical, energy-intensive, or interest-rate sensitive. Their investment projects tend to be large, long-lived and costly to reverse, making them more exposed to higher user costs of capital, tighter credit conditions, weaker cash flow and lower expected demand following a monetary tightening.

The timing of the responses helps distinguish the underlying channels. Construction responds relatively quickly, with significant negative relative effects at $h = 4$ and $h = 8$, consistent with exposure to credit conditions, property-market expectations, mortgage-sensitive demand and project finance. Solid fuels and oil refining shows a slower but more persistent contraction, consistent with long capital-planning cycles and specialised capital stock. Other manufacturing becomes more contractionary by $h = 16$, suggesting that financial-accelerator effects may build gradually as weaker demand and higher financing costs reduce internal funds and investment capacity.

Positive relative responses for Engineering and vehicles and Textiles, clothing, leather and footwear indicate weaker exposure to the domestic interest-rate channel rather than an

expansionary effect of monetary tightening. Engineering and vehicles reach 3.24 percentage points at $h = 16$. As an export-oriented and internationally integrated manufacturing sector, its investment is shaped not only by domestic financing conditions, but also by global demand, exchange-rate movements, supply-chain adjustment, and discrete capital projects. Textiles reach 8.51 percentage points at $h = 12$; however, this sector has a small investment base, meaning that a limited number of lumpy capital expenditures can generate large measured relative responses. These coefficients therefore show that both sectors respond less negatively than the rest of the panel over the relevant horizons.

The main conclusion is therefore that the UK investment channel operates through structurally uneven sectoral channels, with the strongest relative effects concentrated in capital-intensive, cyclical, and finance-sensitive industries. The economic implication is that monetary policy may affect the composition of investment as well as its aggregate level. If tightening disproportionately reduces investment in construction, energy-related sectors, and parts of manufacturing, the effects may extend to future productive capacity, supply-chain resilience, and regional activity. These sectors are important for capital formation and are unevenly distributed across the UK, so heterogeneous investment responses may translate into uneven real-economy effects across industries and locations.

6.1 Policy Implications

The findings are relevant for policymakers assessing how monetary policy transmits to real activity through business investment. The aggregate response is negative but statistically insignificant, yet the sectoral results show significant heterogeneity. This means that aggregate investment data may conceal stronger adjustment in rate-sensitive sectors, including construction, energy-related industries, and parts of manufacturing. This pattern is consistent with firm and sector-level literature discussed previously.

Construction is particularly notable. Its strong negative response, a quarter of a century after Ganley and Salmon (1997), suggests that the structural features generating sectoral heterogeneity are durable rather than period specific. The implication is not that monetary policy should target individual sectors; Bank Rate remains an aggregate instrument directed at price stability. It is rather that sectoral investment indicators warrant monitoring alongside

aggregate measures, since a statistically insignificant aggregate may reflect offsetting sectoral responses rather than weak policy effects.

The policy implication is, therefore, that sectoral composition matters. If tightening disproportionately reduces investment in construction, energy-related sectors, and manufacturing, the effects may extend beyond short-run demand management to medium-term productive capacity, supply-chain resilience, and regional activity.

6.2 Limitations and Future Research

As with all empirical research, the results should be interpreted alongside several limitations. One limitation is the level of aggregation in the data. The analysis uses 20 broad industry groups, but firms within the same sector may differ substantially in leverage, liquidity, size, export exposure, and reliance on external finance. This means that some firm-level transmission channels may be hidden within the sectoral estimates. Further research could combine high-frequency monetary policy shocks with microdata such as sector-level measures of capital intensity, energy intensity, leverage, liquidity or external finance dependence, to identify which characteristics explain the strongest responses.

Further, this dissertation focuses on the Target factor. This is appropriate for isolating surprises to the current policy rate, but it does not capture the full term-structure dimension of monetary policy. Further research could extend the framework to information-adjusted Path shocks to examine whether forward guidance and expected future rates generate different sectoral investment responses.

Finally, future research should trace these heterogeneous investment responses through to output, employment, productivity, and regional activity, closing the transmission feedback loop from monetary policy through sectoral capital formation to aggregate real outcomes.

7. References

- Abadie, A., Athey, S., Imbens, G.W. and Wooldridge, J.M. (2023) 'When should you adjust standard errors for clustering?', *The Quarterly Journal of Economics*, 138(1), pp. 1–35.
- Bernanke, B.S. and Blinder, A.S. (1988) 'Credit, money, and aggregate demand', *American Economic Review*, 78(2), pp. 435–439.
- Bernanke, B.S. and Blinder, A.S. (1992) 'The federal funds rate and the channels of monetary transmission', *American Economic Review*, 82(4), pp. 901–921.
- Bernanke, B.S. and Gertler, M. (1989) 'Agency costs, net worth, and business fluctuations', *American Economic Review*, 79(1), pp. 14–31.
- Bernanke, B.S., Gertler, M. and Gilchrist, S. (1999) 'The financial accelerator in a quantitative business cycle framework', in Taylor, J.B. and Woodford, M. (eds.) *Handbook of Macroeconomics*, Vol. 1C. Amsterdam: Elsevier, pp. 1341–1393.
- Bernanke, B.S. and Gertler, M. (1995) 'Inside the black box: the credit channel of monetary policy transmission', *Journal of Economic Perspectives*, 9(4), pp. 27–48.
- Bunn, P., Oikonomou, M., Anayi, L., Mizen, P., Thwaites, G. and Bloom, N. (2021) 'Influences on investment by UK businesses: evidence from the Decision Maker Panel', *Bank of England Quarterly Bulletin*, 2021 Q2. Available at: <https://www.bankofengland.co.uk/quarterly-bulletin/2021/2021-q2/influences-on-investment-by-uk-businesses-evidence-from-the-decision-maker-panel>
- Burr, N. and Willems, T. (2024) 'About a rate of (general) interest: how monetary policy transmits', *Bank of England Quarterly Bulletin*, 2024. Available at: <https://www.bankofengland.co.uk/quarterly-bulletin/2024/2024/about-a-rate-of-general-interest-how-monetary-policy-transmits>
- Cameron, A.C., Gelbach, J.B. and Miller, D.L. (2011) 'Robust inference with multiway clustering', *Journal of Business & Economic Statistics*, 29(2), pp. 238–249.
- Cameron, A.C. and Miller, D.L. (2015) 'A practitioner's guide to cluster-robust inference', *Journal of Human Resources*, 50(2), pp. 317–372.
- Cesa-Bianchi, A., Thwaites, G. and Vicondoa, A. (2020) 'Monetary policy transmission in the United Kingdom: a high frequency identification approach', *European Economic Review*, 123, 103375.
- Choi, S., Willems, T. and Yoo, S.Y. (2024) 'Revisiting the monetary transmission mechanism through an industry-level differential approach', *Journal of Monetary Economics*.
- Christiano, L.J., Eichenbaum, M. and Evans, C.L. (1999) 'Monetary policy shocks: what have we learned and to what end?', in Taylor, J.B. and Woodford, M. (eds.) *Handbook of Macroeconomics*, vol. 1A. Amsterdam: Elsevier, pp. 65–148.
- Christiano, L.J., Eichenbaum, M. and Evans, C.L. (2005) 'Nominal rigidities and the dynamic effects of a shock to monetary policy', *Journal of Political Economy*, 113(1), pp. 1–45.
- Cloyne, J. and Hürtgen, P. (2016) 'The macroeconomic effects of monetary policy: a new measure for the United Kingdom', *American Economic Journal: Macroeconomics*, 8(4), pp. 75–102.
- Cloyne, J., Ferreira, C., Froemel, M. and Surico, P. (2023) 'Monetary policy, corporate finance, and investment', *Journal of the European Economic Association*, 21(6), pp. 2586–2634.

- Dedola, L. and Lippi, F. (2005) 'The monetary transmission mechanism: evidence from the industries of five OECD countries', *European Economic Review*, 49(6), pp. 1543–1569.
- Dixit, A.K. and Pindyck, R.S. (1994) *Investment under uncertainty*. Princeton, NJ: Princeton University Press.
- Fazzari, S.M., Hubbard, R.G. and Petersen, B.C. (1988) 'Financing constraints and corporate investment', *Brookings Papers on Economic Activity*, 1988(1), pp. 141–206.
- Ganley, J. and Salmon, C. (1997) *The industrial impact of monetary policy shocks: some stylised facts*. Bank of England Working Paper No. 68. London: Bank of England.
- Gertler, M. and Gilchrist, S. (1994) 'Monetary policy, business cycles, and the behavior of small manufacturing firms', *The Quarterly Journal of Economics*, 109(2), pp. 309–340.
- Gertler, M. and Karadi, P. (2015) 'Monetary policy surprises, credit costs, and economic activity', *American Economic Journal: Macroeconomics*, 7(1), pp. 44–76.
- Gürkaynak, 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), pp. 55–93.
- Hall, R.E. and Jorgenson, D.W. (1967) 'Tax policy and investment behavior', *American Economic Review*, 57(3), pp. 391–414.
- Hill, R.C., Griffiths, W.E. and Lim, G.C. (2008) *Principles of Econometrics*. 3rd edn. Hoboken, NJ: Wiley.
- Jarociński, M. and Karadi, P. (2020) 'Deconstructing monetary policy surprises — the role of information shocks', *American Economic Journal: Macroeconomics*, 12(2), pp. 1–43.
- Jorgenson, D.W. (1963) 'Capital theory and investment behavior', *American Economic Review*, 53(2), pp. 247–259.
- Kashyap, A.K. and Stein, J.C. (2000) 'What do a million observations on banks say about the transmission of monetary policy?', *American Economic Review*, 90(3), pp. 407–428.
- Kuttner, K.N. (2001) 'Monetary policy surprises and interest rates: evidence from the Fed funds futures market', *Journal of Monetary Economics*, 47(3), pp. 523–544.
- Miranda-Agrippino, S. and Ricco, G. (2021) 'The transmission of monetary policy shocks', *American Economic Journal: Macroeconomics*, 13(3), pp. 74–107.
- Peersman, G. and Smets, F. (2003) 'The monetary transmission mechanism in the euro area: evidence from VAR analysis', in Angeloni, I., Kashyap, and Mojon, B. (eds.) *Monetary policy transmission in the euro area*. Cambridge: Cambridge University Press, pp. 36–55.
- Peersman, G. and Smets, F. (2005) 'The industry effects of monetary policy in the euro area', *The Economic Journal*, 115(503), pp. 319–342.
- Plagborg-Møller, M. and Wolf, C.K. (2021) 'Local projections and VARs estimate the same impulse responses', *Econometrica*, 89(2), pp. 955–980.
- Ramey, V.A. (2016) 'Macroeconomic shocks and their propagation', in Taylor, J.B. and Uhlig, H. (eds.) *Handbook of Macroeconomics*, Vol. 2. Amsterdam: Elsevier, pp. 71–162.
- Romer, C.D. and Romer, D.H. (2004) 'A new measure of monetary shocks: derivation and implications', *American Economic Review*, 94(4), pp. 1055–1084.
- Sims, C.A. (1980) 'Macroeconomics and reality', *Econometrica*, 48(1), pp. 1–48.

Sims, C.A. (1992) 'Interpreting the macroeconomic time series facts: the effects of monetary policy', *European Economic Review*, 36(5), pp. 975–1000.

Swanson, E.T. (2021) 'Measuring the effects of Federal Reserve forward guidance and asset purchases on financial markets', *Journal of Monetary Economics*, 118, pp. 32–53.

Tobin (1969) 'A general equilibrium approach to monetary theory', *Journal of Money, Credit and Banking*, 1(1), pp. 15-29.