Title

Essays on Monetary Policy and Financial Stability

Author

Uppal, Ali

Publication Date

2024

UNIVERSITY OF CALIFORNIA SAN DIEGO

Essays on Monetary Policy and Financial Stability

A dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Economics

by

Ali Uppal

Committee in charge:

Professor James Hamilton, Chair
Professor Juan Herreño
Professor Nir Jaimovich
Professor Allan Timmermann
Professor Fabian Trottner
Professor Johannes Wieland

2024

The Dissertation of Ali Uppal is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2024

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

I am forever indebted to the numerous people who have supported and encouraged me throughout this journey.

I am immensely grateful to my advisor, Jim Hamilton, for his wisdom, guidance, and importantly, his belief in me, not only for the duration of the PhD, but from our very first interaction during the PhD visit day.

I have also benefited greatly from my other committee members: Juan Herreño, Nir Jaimovich, Allan Timmermann, Fabian Trottner, and Johannes Wieland. Their insightful feedback and thought-provoking questions have both enriched my research and challenged me to be a better economist.

I would like to thank my friends at UCSD: Anindo Sarkar, Beata Łuczywek, Cole Dreier, Meenakshi Balakrishna, Molly Shatto, and Steven Yee. They not not only made this journey much more fun but were a consistent source of helpful advice. A special shout-out to Molly and Steven for being indispensable to the development of my job market paper as well as being a never-ending second pair of eyes, and to Beata for her camaraderie during the job market process.

I would also like to express my appreciation for Mike Boozer and Andrew Metrick at Yale University. Their mentorship, guidance, and friendship were instrumental to my academic development and pursuit of a doctorate.

At a personal level, I would like to express my sincerest appreciation for my family and friends in London. Thank you to my mum, Saeeda Batool, for her unfaltering belief in me and for making every effort and sacrifice in order to see me succeed. Her unconditional love and support have been invaluable not only during this period, but throughout my life. Thank you to my sister, Zena Uppal, for taking on numerous responsibilities so effortlessly so that I could embark on this academic endeavour. Thank you to my friends Chih Yen Ku, Minar Miah, Mustafa Bux, Samip Naik, and Sandeep Nijjar; our countless conversations have truly been a great source of joy and motivation.

Finally, I could not have done this journey without my wife, Stéphanie Machabée. She inspired my pursuit of a doctorate and has been my rock throughout this experience. I am forever grateful for her willingness and patience to read everything I write and listen to everything I present. Her balance of discipline and freedom empowered me to complete this dissertation while also making these years some of the most enjoyable in my life. I would be remiss if I did not mention the birth of our son, Elyas Shahzad Macpal, during the job market in the final year of my PhD. His presence granted me perspective during what is normally an especially stressful time and I am so grateful to have him see me graduate.

Chapter 1 is currently being prepared for submission for publication of the material. The dissertation author was the sole investigator and author of this paper.

Chapter 2 is currently being prepared for submission for publication of the material. The dissertation author was the sole investigator and author of this paper.

Chapter 3 is currently being prepared for submission for publication of the material. Giampaolo Bonomi and the dissertation author, Ali Uppal, are co-authors of this chapter.

Chapter 4 is currently being planned for submission for publication of the material. Fabian Trottner and the dissertation author, Ali Uppal, are co-authors of this chapter.

VITA

2012	Bachelor of Science, London School of Economics
2012–2015	Economist, HM Treasury
2015	Economist, Bank of England
2016	Master of Arts, Yale University
2016	Economist, International Monetary Fund
2016–2018	Head of Balance Sheet Analysis, HM Treasury
2018	Fiscal Expert, International Monetary Fund
2018–2019	Researcher, London Business School
2020–2022	Research Assistant, University of California San Diego
2022–2023	NBER Rapporteur, National Bureau of Economic Research
2023	Dissertation Fellow, Federal Reserve Board
2020–2024	Teaching Assistant, University of California San Diego
2024	Doctor of Philosophy, University of California San Diego

ABSTRACT OF THE DISSERTATION

Essays on Monetary Policy and Financial Stability

by

Ali Uppal

Doctor of Philosophy in Economics

University of California San Diego, 2024

Professor James Hamilton, Chair

This dissertation consists of four chapters, each of which studies monetary policy, financial stability, or their interaction.

Chapter one shows empirically that, contrary to theoretical claims, raising interest rates *increases* bank leverage. I propose and empirically validate the loan-loss mechanism to explain this result: contractionary shocks increase loan losses, reduce bank profits and equity, and ultimately increase bank leverage. I develop a banking model where floating-rate loans entail a trade-off between interest rate risk and credit risk, which generates the loan-loss mechanism. Using microdata, I provide empirical evidence consistent with floating-rate loans hedging interest rate risk at the expense of generating loan losses.

Chapter two examines the effects of central bank meetings on stock returns. Cieslak et al. (2019) show that stock returns in the US and internationally are driven by even-week meetings of the Federal Open Market Committee. I find that the US result and the proposed mechanism do not hold out-of-sample, losing robustness as early as 2004. Prior to 2004, there appear to be outliers driving the result. Finally, I show that the international result does not apply in either the UK or Japan.

Chapter three studies the consequences of including financial stability among the central bank's objectives when market players are strategic. Our model predicts that central banks underreact to economic shocks, a prediction consistent with the Federal Reserve's behaviour during the 2023 banking crisis. Policymakers' stability concerns bias investors' choices, inducing inefficiency. If central banks have private information about their policy intentions, the equilibrium forward guidance is vague because fully informative communication is not credible. A "kitish" central banker, who is less concerned about stability, reduces these inefficiencies.

Chapter four studies how financial and production networks affect the transmission of financial shocks to the real economy. We propose a general equilibrium model of production networks featuring heterogeneous banks and endogenous firm-bank linkages. We theoretically characterise the aggregate effects of bank-specific shocks in terms of a number of sufficient statistics. We suggest an approach to empirically complement our theoretical framework which relies on misconduct provisions of UK banks combined with detailed firm-bank-loan data to construct instruments for firms' credit supply.

Chapter 1

Do Higher Interest Rates Make The Banking System Safer? Evidence From Bank Leverage

1.1 Introduction

As early as 1945, one of the most influential economists of the twentieth century, Paul Samuelson, proclaimed that in response to an increase in interest rates, the banking system is "tremendously better off" (Samuelson (1945)).[1] Close to seventy years later, former Federal Reserve Chair Janet Yellen echoed a similar sentiment in an influential speech on monetary policy and financial stability, explaining that higher interest rates reduce financial sector vulnerabilities through reducing their leverage (i.e., ratio of assets to equity).[2] The empirical validity of these claims is critical to improving our understanding of the transmission of monetary policy through banks in addition to informing the ongoing debate on whether monetary policy should support financial stability. In this paper, I explicitly address this first-order question: do contractionary monetary policy shocks actually improve bank vulnerability by reducing bank leverage?

The answer to this question, from much of the theoretical literature, is yes. Woodford (2012) argues, using a typical New Keynesian model with credit frictions, that "it is appropriate

[1] Specifically, Samuelson (1945) argues that an interest rate hike would significantly improve the profitability and stability of the banking system.

[2] https://www.federalreserve.gov/newsevents/speech/yellen20140702a.htm

to use monetary policy to 'lean against' a credit boom" which in his model implies tightening monetary policy to reduce leverage. Angeloni and Faia (2013) build a dynamic macroeconomic model featuring banks to similarly conclude that "the increase in interest rate activates the risk taking channel: bank leverage and risk decline." Dell'Ariccia et al. (2014) develop a model of financial intermediation where banks engage in costly monitoring to reduce the credit risk in their loan portfolios. Despite their different modelling approach, they reach the same conclusion: "a reduction in risk-free interest rates leads banks to increase their leverage" where the risk-free rate refers to the policy rate. Drechsler et al. (2018b) take yet another approach by developing a dynamic asset pricing model in which monetary policy affects the risk premium component of the cost of capital. Nonetheless, their analysis leads to the same claim: "Lower nominal rates make liquidity cheaper and raise leverage." Martinez-Miera and Repullo (2019) extend the banking model of Martinez-Miera and Repullo (2017) to include monetary policy and similarly show that "[monetary] tightening reduces aggregate investment. . . and reduces bank leverage." Finally, Martin et al. (2021) highlight that the framework of Ghote (2021), which consists of monetary and macroprudential policy intervention in a general equilibrium economy with recurrent boom-bust cycles, also supports leaning against a boom.[3] In particular, Martin et al. (2021) summarise the theoretical literature by concluding the following: "This is true in most models. . . By tightening ex ante, monetary policy contributes to reducing credit and, more specifically, leverage."[4]

Given such strong and consistent claims across much of the theoretical literature and a plethora of modelling approaches, one might expect considerable empirical support. However, as Boyarchenko et al. (2022) highlight in their review paper, there is very limited and inconclusive empirical evidence on the causal impact of monetary policy on leverage and no empirical

[3] Macroprudential policy uses primarily regulatory measures (e.g., bank leverage requirements) to limit financial crisis risk. See for example Galati and Moessner (2018) and Gourio et al. (2018).

[4] See Appendix A.1 for a brief summary of these models and the underlying mechanisms.

evidence of the underlying mechanism.[5,6]

The first contribution of this paper is to provide robust empirical evidence of the impact of contractionary monetary policy shocks on bank leverage. My empirical strategy relies on existing measures of exogenous monetary policy shocks which capture unexpected changes in the Fed Funds Rate (FFR).[7] Using quarterly data from the Federal Deposit Insurance Corporation between 1984 and 2007, I estimate lag-augmented local projections of aggregate bank leverage with these exogenous shocks. My first finding is contrary to much of the theoretical literature. I find that a contractionary monetary shock that induces a one percentage point rise in the FFR leads to a five to ten percent *increase* in bank leverage. I show that this finding is robust to different definitions of leverage, different time-periods, different lag lengths, and different monetary policy shock series.

My second contribution is to document empirically a mechanism that can explain why leverage increases in response to contractionary monetary policy shocks. This is important as it sheds light on how banks respond to and are affected by monetary policy. The literature has traditionally focused on the bank lending channel as the main way in which banks interact with monetary policy whereby contractionary monetary policy reduces bank lending (see for example Kashyap and Stein (1994) and Bernanke and Gertler (1995)). However, the last decade has seen a resurgence of research on the transmission of monetary policy through the financial system, largely driven by empirical evidence that monetary policy has meaningful consequences on financial institutions in ways that are not captured by the workhorse New Keynesian models (Drechsler et al. (2018a)). I show that while raising interest rates does indeed reduce bank

[5] The main empirical papers that are related to this question include Miranda-Agrippino and Rey (2020), Wieland and Yang (2020), and Li (2022). However, the primary focus of these papers is not the estimation of a domestic bank leverage response to domestic monetary policy shocks. The inclusion of such estimates in these papers are part of ancillary analyses and, as such, the analyses are not supported with sufficient robustness checks nor detailed discussion of potential mechanisms.

[6] While not directly examining the impact of monetary policy on leverage, Grimm et al. (2023) examine the impact of loose monetary policy on financial instability. Specifically, they find that extended periods of accomodative policy, when followed by a tightening, can increase the likelihood of financial distress.

[7] I use several different measures of exogenous monetary policy shocks including Romer and Romer (2004), Gertler and Karadi (2015), and Bu et al. (2021).

borrowing (as per the bank lending channel), it also increases the proportion of loans that are delinquent and so increases loan losses. Unexpectedly higher loan losses decrease bank profits which subsequently reduce bank equity.[8] I find that the drop in equity increases leverage more than the drop in borrowing decreases leverage and so bank leverage increases overall. I term this the *loan-loss mechanism*. Moreover, instead of relying on ruling out alternative mechanisms to evaluate the importance of the loan-loss mechanism, I take advantage of accounting identities which allow me to show precisely that the loan-loss mechanism explains almost all the variation in bank leverage in response to contractionary monetary policy shocks. Finally, I show that while loan losses and leverage increase in response to monetary policy shocks (where the FFR rises), loan losses do not rise and leverage actually falls in response to contractionary oil shocks (where the FFR does not rise). This analysis provides suggestive evidence, at the aggregate level, of the importance of the rise in the FFR and hence the potential role of floating-rate loans in generating loan losses.

My next contribution is to dissect the theoretical literature in order to show where and why so many, and such different, models generate empirically inconsistent leverage responses. Investigating the literature in this way is important not only to provide an empirically-grounded theoretical answer to whether contractionary monetary policy reduces bank leverage, but also because bank leverage, per se, plays a vital role in macroeconomic models with financial sectors. For example, as highlighted in Adrian et al. (2014), in many models, such as Fostel and Geanakoplos (2008) and Geanakoplos (2010), when the bank's own funds are fixed, leverage is the key state variable and lending is determined solely by leverage. This directly connects leverage to the bank lending channel of monetary policy. Furthermore, as commented in Ajello et al. (2022), leverage is core to the financial accelerator models (e.g., Bernanke et al. (1999)), and typically both amplifies and propagates the response of the economy to shocks, thus generating aggregate fluctuations.

[8] English et al. (2018) also find that contractionary monetary policy reduces bank profits while Altavilla et al. (2018) find that a prolonged period of low interest rates reduces loan losses which boosts bank profits.

I show that the aforementioned empirical inconsistency of the literature appears to derive from three broad, though not necessarily mutually exclusive, modelling decisions. First are models such as Angeloni and Faia (2013) and Drechsler et al. (2018b) that rely on some form of a substitution effect as the dominant mechanism through which higher interest rates reduce bank leverage. Second are models such as Woodford (2012) and Rannenberg (2016) that incorrectly rely on the observed procyclical behaviour of leverage in order to conclude that leverage declines in response to monetary policy tightening. Finally, models such as Gertler and Kiyotaki (2010) and Gertler and Karadi (2011), while able to generate an increase in leverage in response to contractionary monetary policy shocks, attribute this rise to an increase in expected profitability. However, despite an empirically consistent leverage response, the proposed mechanism is inconsistent with the observed evidence that profitability falls rather than rises in response to a monetary contraction. Moreover, this class of models typically generates an increase in leverage in response to any contractionary shock which is inconsistent with the empirical response of leverage to other (non-monetary) contractionary shocks, such as oil shocks. This underscores the importance of the underlying loan-loss mechanism which is specific to contractionary monetary policy shocks.

The empirical inconsistencies across models typically arise in the the banking block rather than the general equilibrium structure of the model. Indeed, some banking-specific models generate more empirically consistent dynamics as they feature both a fall in profits and a rise in leverage in response to contractionary monetary policy shocks.[9] However, such models still do not capture a loan-loss mechanism which features a role for floating-rate loans. Therefore, I develop a banking model, building on Kirti (2020), that emphasises the role of floating-rate loans and credit risk. In my model, banks optimise by choosing the floating share of their loan portfolio. While floating-rate loans hedge against interest rate risk, they do so by passing this

[9] For example, Van den Heuvel (2009) develops a bank capital channel of monetary policy which sees profits fall and leverage rise following contractionary monetary policy due to maturity transformation while Corbae and Levine (2024) also see profits fall and leverage rise in response to contractionary monetary policy as higher rates induce greater risk-taking.

risk onto borrowers which generates credit risk for the bank. As such, a key insight of the model is that banks are doing risk transformation, and that this implies a trade-off between managing interest rate risk and credit risk. The model generates implications for the data that depend on the share of a bank's loan portfolio that is floating rate. Specifically, the model predicts that banks with a higher share of floating-rate loans will see greater loan losses in response to a contractionary monetary policy shock.

Finally, I use microdata, in particular bank-level variation in the share of floating-rate loans, in a panel local projection framework to test the implications of the model. Consistent with the model, I find that banks with higher shares experience higher net interest income but also higher loan losses in response to contractionary monetary policy shocks. The effect on profits is ultimately negative for those with higher shares which results in a larger increase in leverage. This provides further evidence for the role of floating-rate loans in generating the loan-loss mechanism and has important implications for regions where loans are predominantly floating-rate (e.g., Europe).

The overall contributions of this paper lend support to the conclusions of Svensson (2017), Svensson (2018), and former Federal Reserve Chair Ben Bernanke that monetary policy should not target financial stability while also documenting an important channel through which tighter monetary policy adversely affects the stability of the banking system.[10] Svensson argues that an important cost of monetary policy that targets financial stability is that it weakens the economy by allowing lower inflation and/or higher unemployment than would otherwise be the case, hence reducing the economy's resilience to shocks. Importantly, these papers argue that this cost is larger than the potential benefit of reducing the probability of a crisis (e.g., by reducing bank leverage). However, even in these papers, there is an implicit acceptance that tight monetary policy reduces bank leverage. Without that key benefit, the argument in favour of monetary policy targeting financial stability appears substantially weaker, and the conclusions of Svensson (2017) and Svensson (2018) significantly stronger. Bernanke, Svensson, and I end up

[10] https://www.brookings.edu/articles/should-monetary-policy-take-into-account-risks-to-financial-stability/

concurring with the Tinbergen (1952) rule which asserts that we need at least *n* policy instruments for *n* policy goals. Therefore, given the significant adverse impact contractionary monetary policy has on leverage, especially when the share of floating-rate loans is high, monetary policy should focus on its traditional mandate of price stability, leaving issues of financial stability to macroprudential policy.

The remainder of this paper is as follows. Section 1.2 describes the data used. Section 1.3 documents the time-series evidence. Section 1.4 explains where and why the theoretical literature is empirically inconsistent while Section 1.5 develops a model to highlight the role of floating-rate loans. Section 1.6 explores the implications of the model using microdata. Section 1.7 provides a conclusion. Appendix A.1 summarises theoretical papers that predict contractionary monetary shocks reduce bank leverage, Appendix A.2 provides further discussion about the differences between book leverage and market leverage, Appendix A.3 presents robustness checks in relation to the time-series evidence, and Appendix A.4 provides details of the theoretical model.

1.2 Data

This paper brings together both aggregate and individual bank-level data, ensuring that the individual bank-level data matches the aggregate data, and combines it with the Fed Funds Rates and different measures of monetary policy shocks. All data is either already at quarterly frequency or has been transformed to be at quarterly frequency. Finally, the data coverage is from the first quarter of 1984 up until the last quarter of 2006. Given that the 2007-08 global financial crisis (GFC) resulted in such substantive changes to the regulatory architecture, my analysis will focus on the period prior to the crisis. This section describes the different data and their sources.

1.2.1 Aggregate Banking Sector Time-Series Data

The aggregate banking sector time-series data is from the Federal Deposit Insurance Corporation (FDIC). Specifically, I obtain aggregated balance sheet and income statement data for all FDIC-insured institutions for each quarter starting in 1984 using the FDIC's Quarterly Banking Profile data. This provides me with accounting-based measures of different variables.[11]

From the aggregate balance sheet, I collect time-series data on four key variables. The first variable is total banking sector assets. While useful in its own right, the measure of total assets will mostly be used to normalise all remaining variables so that they are interpreted as a share of total assets. Second, I collect data on loans that are 30-89 days past due. This simply measures loans where the borrower is up to three months behind on a payment. The final two variables capture different measures of bank equity. The first is total equity which is also sometimes referred to as the net worth of the bank. Using this measure of equity, we can define the simple leverage of the bank as total assets divided by total equity. This is easily comparable across time, space, and banks. As such, when referring to leverage, I will be referencing simple leverage, unless otherwise specified. The second variable is regulatory equity which is also known as Tier 1 capital. This variable is a stricter definition of equity as it excludes several components from total equity such as revaluation reserves and hybrid capital instruments. The regulatory community argues that the Tier 1 Leverage Ratio (i.e., regulatory equity divided by average assets over the quarter) represents a more accurate measure of the losses a bank can withstand in response to a shock. Therefore, I will also use this measure of regulatory leverage (i.e., total assets divided by regulatory equity) for robustness.[12]

From aggregate income statements, I have four main variables. First, I collect data on aggregate profits as measured by net income from the income statements. Second, I collect data on dividends. The last two variables represent loan losses. Specifically, these variables are loan-loss provisions and net charge-offs. The former captures a bank's expectation of future loan

[11] The accounting-based data is book leverage rather than market leverage. A number of papers (e.g., Adrian et al. (2019)) argue that this is the relevant measure for bank balance sheet decisions. See Appendix A.2 for further discussion as well as a Figure A.6 for a robustness check based on a measure of market leverage from He et al. (2017). The robustness check shows the results are qualitatively similar across market and book leverage measures.

[12] Note that this differs in two minor ways from the regulatory measure used in practice. First, the regulatory measure uses average assets over the quarter (see footnote 5 of https://www.kansascityfed.org/documents/8087/BankCapitalAnalysisTable_December31_2020.pdf) while, for data availability reasons, I use total assets at the end of the quarter. Second, I focus on leverage instead of the leverage ratio (one is just the reciprocal of the other) as it is far more intuitive.

losses, while the latter are recorded when a bank decides to finally write off a loan. If loan-loss provisions were perfectly estimated by banks, they would be exactly equal to net charge-offs over the long run. In the 10 years prior to the financial crisis, loan-loss provisions averaged around 110% of net charge-offs, which is consistent with regulatory examiners pushing for conservative estimates of expected losses.[13] Therefore, while loan-loss provisions might be slightly conservatively estimated, my analysis utilises provisions instead of net charge-offs for two reasons. First, provisions are recognised in a timelier fashion than charge-offs. Indeed, as soon as a shock occurs, banks will update their estimate of expected loss in accordance with accounting standards. Second, provisions directly impact bank profits and subsequently bank equity so there is a direct accounting-identity link between provisions and bank leverage, which will be important for my empirical work.[14] Nonetheless, the underlying mechanism in my empirical analysis remains the same whether one uses provisions or net charge-offs as both follow a very similar pattern in response to a contractionary monetary policy shock (see Figure A.1 in Appendix A.3).

1.2.2 Bank-Level Panel Data

Bank-level data requires a consistent time series that is merger adjusted. Therefore, I use the bank-level series of Drechsler et al. (2017). My bank-level analysis uses all the variables in the aggregate data and also includes the share of the loan portfolio that is effectively floating-rate. This floating-rate data starts from the second quarter of 1997. As such, the panel data analysis will all be after 1999 (to allow for sufficient lags). In Section 1.6, I show that the underlying bank-level data very closely matches the aggregate data for my sample of 1997 to 2006.

[13] https://fraser.stlouisfed.org/title/economic-trends-federal-reserve-bank-cleveland-3952/economic-trends-november-5-2015-529746/loan-loss-provisioning-517772

[14] Note that there are some concerns that banks may manipulate the timing of loan-loss provisions for tax advantages. However, the 1969 and 1986 Tax Reform Acts largely removed these incentives (see Walter (1991)).

1.2.3 Monetary Policy Data

The monetary policy data has two components. The first is simply the Fed Funds Rate (FFR) which is directly from FRED. The second set of monetary policy data is more substantive. Specifically, I collect a number of different estimates of exogenous changes in monetary policy (i.e., monetary policy shocks). There is a large literature on constructing monetary policy shocks and a number of papers that compare and contrast the different shocks (see for example Ramey (2016)). This paper does not seek to evaluate the effectiveness of a given monetary policy shock measure. Instead, it focuses on how bank leverage responds to a given exogenous monetary policy shock. A benchmark monetary policy shock used in the literature is the shock series by Romer and Romer (2004) (hereafter the RR shock). Their identification strategy combined narrative methods with the Federal Reserve's (the Fed) own internal forecasts. Specifically, they used narrative methods to deduce a series of intended changes to the interest rate during the Fed's monetary policy meetings. Moreover, they separated the endogenous response of policy to information about the economy from the desired exogenous shock by regressing the intended funds rate change on the current rate and on the internal forecasts. The residuals of this regression are essentially the monetary policy shock. I use the updated RR shock series from Wieland and Yang (2020) which allows me to have a quarterly shock from 1984 to 2006. Given its prominence in the literature, the RR shock will be the monetary policy shock used in my baseline specification.

However, there have been specific concerns with the RR shock series. For example, Coibion (2012) finds that the results of Romer and Romer (2004), based on the RR shock series, are particularly sensitive to including the time period 1979-1982 as well as the number of lags. The former will not be an issue in my empirical work as my data start in 1984. I show that the latter is not an issue as my results are relatively robust to varying the number of lags.

To ensure my empirical results are not dependent on one specific measure of monetary policy shocks, I repeat my analysis with two additional monetary policy shock series that have a sufficient time series. I also choose shock series that are estimated using different identification strategies and as such have different features. While the RR shock relies on narrative identification, Gertler and Karadi (2015) (GK shock series) rely on high frequency identification, and Bu et al. (2021) (BRW shock series) utilise a heteroskedasticity-based partial least squares approach, combined with Fama-MacBeth style cross-sectional regressions. However, unlike the RR shock series which covers my entire sample, the GK shock series starts in 1990 and the BRW shock series does not start until 1994. One important feature of the BRW shock series is that Bu et al. (2021) show it contains no significant information effect.

1.2.4 Non-Monetary Shocks Data

To better understand the loan-loss mechanism, I utilise two additional series. While a contractionary monetary policy shock features a rise in the FFR and a decline in GDP, the two additional series have different economic implications which allows me to disentangle the drivers of loan losses.

The first series is an oil shock series. Oil shocks behave as a 'cost-push' shock and so typically do not feature a meaningful rise in the FFR but still result in a decline in GDP. The oil shock series I use is from Känzig (2021). He exploits the institutional features of the Organization of the Petroleum Exporting Countries (OPEC) and high-frequency variation in oil futures prices around OPEC announcements to identify an oil supply news shock. The time period of the shock series is sufficient to span my empirical exercise (i.e., 1984-2007).

The second series is not a shock per se but rather a measure of risk perceptions. Therefore, while analysis utilising this measure is more predictive in nature, rather than necessarily causal, it still provides important insights. Specifically, the measure of perceived risk I use is from Pflueger et al. (2020) who define it as the price of volatile stocks (PVS_t). It is calculated as the average book-to-market ratio of low-volatility stocks minus the average book-to-market ratio of high-volatility stocks. Given this definition, when PVS_t is high, agents are optimistic about the economy (e.g., banks report that they are loosening lending standards). Intuitively, one can think of an increase in PVS_t as acting like a positive demand shock and so should result in a rise in the FFR and GDP. I choose this particular measure of risk perceptions for several reasons. First, and perhaps most importantly, Pflueger et al. (2020) introduce this measure to explicitly evaluate risk-centric theories of business cycles (e.g., Caballero and Simsek (2020b)). Such theories present one important avenue for understanding the interactions between monetary policy and financial stability. Indeed, such models have been extended to show how monetary policy that leans against the wind might have financial stability benefits (see Caballero and Simsek (2020a)).[15] Second, changes in this financial market measure forecast well changes in the real economy. Finally, the time period of the series spans my empirical exercise.

1.3 Time-Series Evidence

My overall empirical approach uses existing measures of exogenous monetary policy shocks in the Jordà (2005) local projection method to estimate impulse responses using data from the start of 1984 until the end of 2006 (unless otherwise specified). This is sometimes referred to as the LP-IV approach (see for example Stock and Watson (2018)). Specifically, I estimate the following for each variable z at each horizon h:

$$z_{t+h} = \alpha_h + \sum_{l=0}^{L} \beta_{h,l} Shock_{t-l} + \sum_{m=1}^{M} \gamma_{h,m} z_{t-m} + \sum_{q=2}^{4} \delta_q Quarter_{qt} + \varepsilon_{t+h}, \quad h = 0, \ldots, 16 \quad (1.1)$$

where z refers to the outcome variable of interest, *Shock* refers to the exogenous monetary policy shock measure, and *Quarter* represents quarterly dummies. The impulse response function is the sequence $\{\beta_{h,0}\}_{h=0}^{H}$ which captures the response of z at time $t+h$ to the shock at time t. In my baseline specification, the lag length is $L = M = 16$ quarters. In line with recent work by Montiel Olea and Plagborg-Møller (2021) on lag-augmented local projections, I use

[15] Note that in a recent paper, Goldberg and López-Salido (2023) extend the framework of Caballero and Simsek (2020b) and show that leaning against the wind may worsen financial stability.

heteroskedasticity-robust standard errors.[16]

The lead-lag exogeneity condition is an important requirement for my specification, and indeed LP-IV approaches more broadly. Stock and Watson (2018) highlight that the main concern is that the shock at time t is correlated with past values of the outcome variable. As such, they suggest a simple test: the shock (i.e., the instrument) should be unforecastable in a regression of the shock at time t on the lags of the outcome variable (z_t in my case). Therefore, I regress the RR shock on 16 lags of leverage and find little evidence of predictability. Specifically, I find that each lag is individually statistically insignificant, the F-statistic when jointly testing all 16 lags also shows statistical insignificance.

1.3.1 The Response of Leverage to a Monetary Policy Shock

My baseline specification is to estimate (1.1) using data from the first quarter of 1984 until the last quarter of 2006 with 16 lags and the RR shock series. Figure 1.1 below depicts the impulse responses of the FFR and leverage.

[16] Montiel Olea and Plagborg-Møller (2021) highlight that with lag-augmented local projections (i.e., where lags of the outcome variable are included as regressors) it is preferable to use heteroskedasticity-robust standard errors instead of Newey-West. They also explain that in the autoregressive literature, “lag augmentation” refers to the practice of using more lags than suggested by the true autoregressive model.

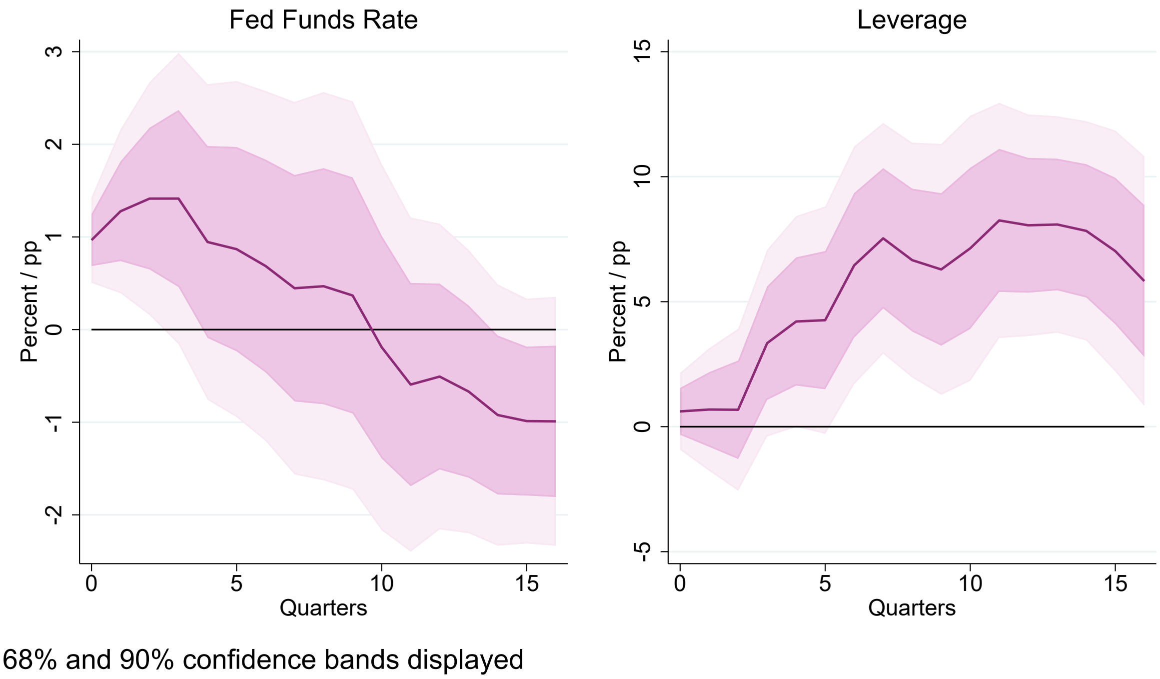

Figure 1.1. Impulse Response of Leverage to Contractionary Monetary Shock

The result above shows that a contractionary monetary policy shock that induces an increase in the FFR of about 1 percentage point significantly increases bank leverage by about 5 percent within a year, which then hovers around 8 percent higher for the remaining three years. This is meaningful response, both in size and persistence. As highlighted earlier, this is in strong contrast to the claims from much of the theoretical literature.

Given the result goes against much of the predictions in the literature and that a core objective of this paper is to provide a robust leverage moment to inform macroeconomic models, it is important to test the robustness of this finding. First, the main result in Figure 1.1 uses the simple definition of leverage described in Section 1.2.1 (i.e., total assets divided by total equity). In Figure A.2, I show the same analysis when using the simple measure of leverage (i.e., total assets divided by regulatory equity). The results do not change in any meaningful way. Next, in Figure A.3, I re-estimate (1.1) using different time periods. Specifically, I reduce the time horizon by three years each time so that I estimate over the period 1987-2006, 1990-2006, and

1993-2006. While the period 1989-92 contained a number of regulatory changes relating to bank leverage, the result is remarkably consistent before and after this period. Given the concerns highlighted by Coibion (2012) about the sensitivity to different lag lengths being used, in Figure A.4, I re-estimate (1.1) with 12 lags, 8 lags, and 4 lags (i.e., 3, 2, and 1 year, respectively). While the precision of the estimates varies across specifications, all of them have leverage rising eventually, though the 8-lag specification to a lesser extent.

My final, and perhaps strictest, robustness test is to use completely different shock series, in particular, ones that use distinct identification strategies. To ensure comparability, I estimate all of them using data from 1994 until 2007 as this is the largest overlapping period. Given the shorter time-horizon, I use 4 lags, otherwise the specification is as in (1.1). Figure A.5 shows the results from using the three different shocks series from Romer and Romer (2004), Gertler and Karadi (2015), and Bu et al. (2021), respectively. Remarkably, the result remains reasonably consistent despite using different shocks. I also repeat this exercise with the measure of market leverage from He et al. (2017). Specifically, in Figure A.6, I show how market leverage responds to each different shock and is qualitatively similar to the results for book leverage and quantitatively larger. The robustness of the result warrants further exploration into the possible mechanisms to understand what is driving the increase in leverage in response to contractionary monetary policy shocks.

1.3.2 The Loan-Loss Mechanism

The literature highlights several different mechanisms that might cause an increase in interest rates to decrease leverage. One of the more intuitive reasons is that higher interest rates make debt financing more expensive relative to equity financing for banks. Given banks decrease the size of their balance sheets in response to contractionary shocks, the decrease in total liabilities will be driven more by a fall in debt liabilities than equity. This substitution effect therefore predicts that higher interest rates reduce bank leverage.

For leverage to rise overall, it must be that the fall in equity is more consequential. As such, I posit an additional mechanism, which I term the *loan-loss mechanism*, that might be driving the overall response in leverage (and offsetting the substitution effect). The mechanism is simple and intuitive and is best described in three key steps. First, a rise in interest rates leads to greater difficulty for borrowers to repay loans which leads to an increase in the proportion of loan repayments that are missed. This should result in *(i) an increasing proportion of loans past due* and *(ii) a delayed but increasing proportion in loan-loss provisions*. The latter rises as banks raise their estimates of expected losses due to the unexpected growth in missed loan repayments. Second, (i) and (ii) imply greater loan losses overall and therefore should result in *(iii) decreasing profits*. Finally, given changes in bank equity are largely driven by changes in profits, decreasing profits should lead to *(iv) decreasing bank equity* and if the overall fall in equity is more important than the fall in assets, then we would expect *(v) increasing leverage* as it is just the ratio of assets to equity.

To test the aforementioned mechanism, I estimate my benchmark specification (i.e., (1.1) with the RR shock, 16 lags, and data from 1984-2006) separately for each of the five variables emphasised in the paragraph above. Specifically, each of the variables will be z in (1.1). Figure 1.2 shows the results of this exercise by showing how these variables respond to a contractionary monetary policy shock.

The first panel (top-left) simply reproduces the impulse response function of the FFR and so the remaining analysis can be interpreted as responding to a monetary policy shock that induces the FFR to increase by around one percentage point on impact. The second panel (top-middle) shows that loans that are up to three months past due increase by nearly 0.5 percentage points as a proportion of total assets at their peak. This is a significant rise as the average share of loans past due during 1984-2006 is around 0.8%. This confirms (i). Similarly, the third panel (top-right) shows that provisions as a proportion of total assets also increase, albeit at a slower pace, which confirms (ii). The greater than 0.1 percentage point rise in provisions as a share of assets is also significant as it is roughly double its average by the end of the projection horizon. The fourth panel (bottom-left) shows that profits as a proportion of total assets decrease by around 0.15 percentage points from an average of around 0.22% at around the same time as when provisions rise which confirms (iii). The fifth panel (bottom-middle) shows regulatory equity falls by nearly five percent within two years and continues to fall to nearly a ten percent decline by the end of the horizon which confirms (iv). Finally, the sixth panel (bottom-right) simply reproduces the main finding in Figure (1.1) (i.e., that leverage rises) and thus confirms (v).

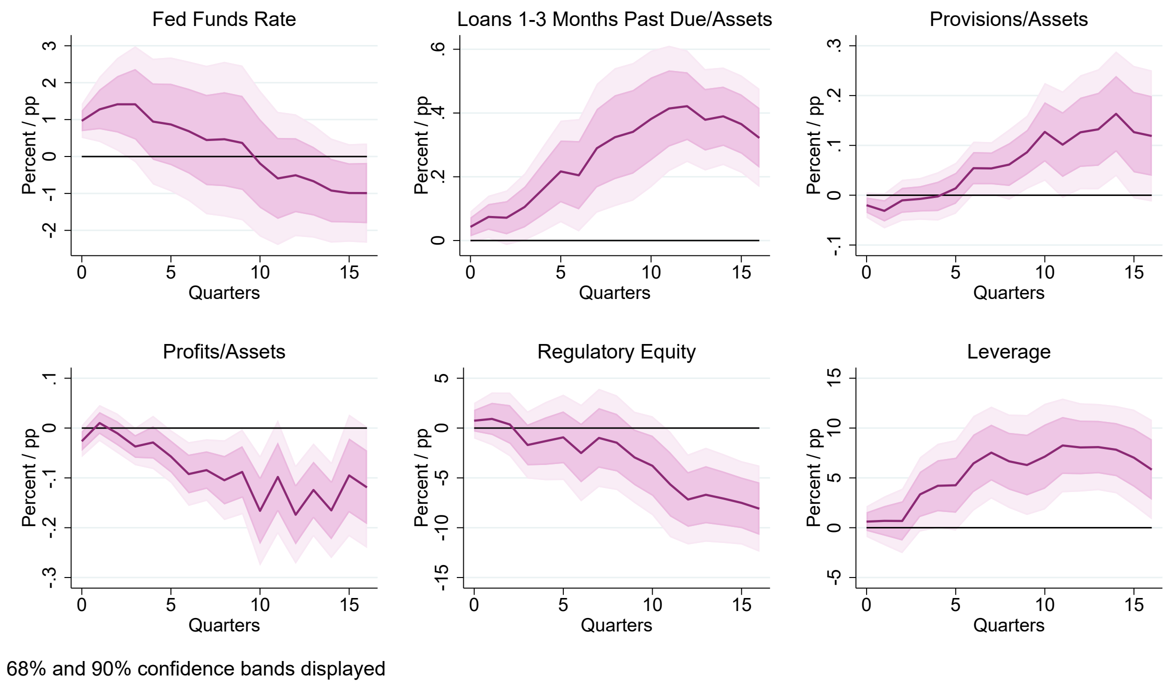

Figure 1.2. Mechanism Underlying Leverage Response

1.3.3 Importance of the Loan-Loss Mechanism

Despite the evidence supporting my proposed mechanism, it is still possible that there are other mechanisms that might be more important in terms of driving the overall increase in leverage. One general approach to deal with this kind of concern is to rule out alternative mechanisms. However, such an approach is not exhaustive as it is difficult to know all possible alternative mechanisms and therefore the best we can usually do is to rule out the most likely contenders. As such, in this section, I take advantage of accounting identities to show that my proposed mechanism explains most of the variation in leverage. Therefore, instead of relying on ruling out possible alternative mechanisms, I show empirically the importance of my mechanism directly.

For my mechanism to be driving the overall response, I need to document two steps. First, that the increase in loan losses, as measured by provisions, in response to the contractionary monetary policy shock (top-right panel of Figure 1.2) is *causing* most, if not all, of the decrease in profits (bottom-left panel of Figure 1.2). Profits can be decomposed into several components on a bank income statement. Specifically, one can utilise the following accounting identity:

$$\frac{\text{Profits (excluding provisions)}_t}{\text{Assets}_t} - \frac{\text{Provisions}_t}{\text{Assets}_t} = \frac{\text{Profits}_t}{\text{Assets}_t} \tag{1.2}$$

where the first term is constructed by adding together net interest income, net noninterest income, net gains on securities, and subtracting taxes. Therefore, if my proposed mechanism is important, it should be the case the variation in profits is driven by the variation in provisions rather than the other income terms. Figure 1.3 below shows the impulse responses of each term in (4.34) which are obtained by estimating (1.1) with each of those terms as the outcome variable z_t.

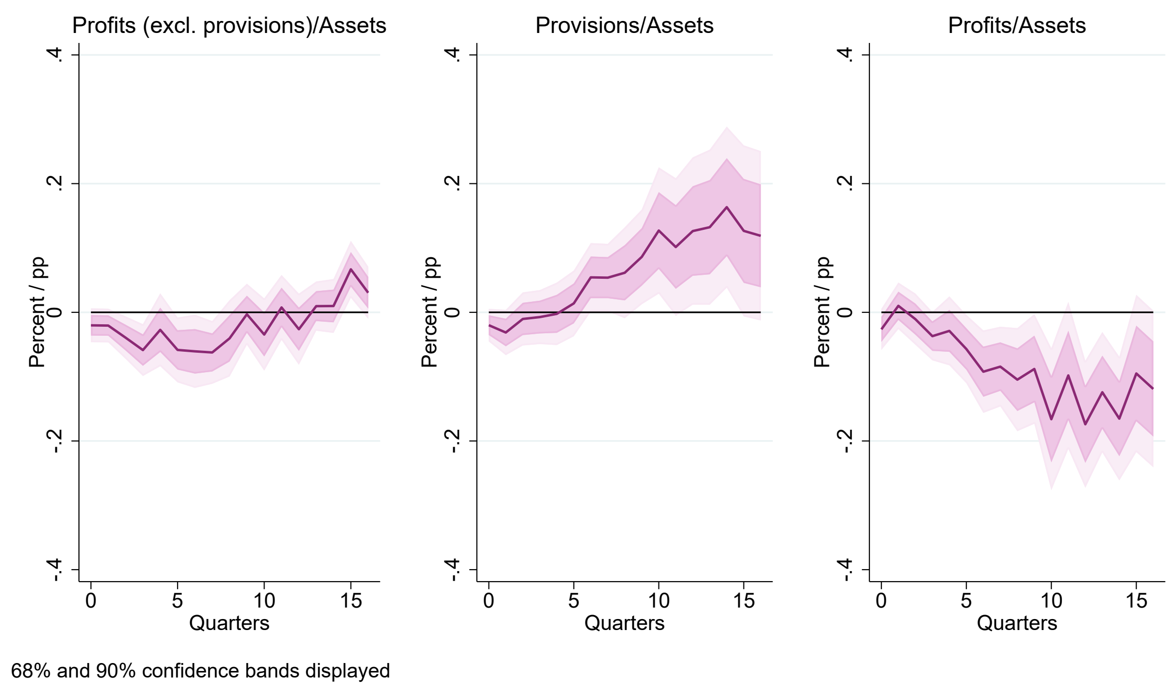

Figure 1.3. Decomposing the Profit Decline

The variation in overall profits is almost entirely driven by the variation in provisions with the remaining variation (captured by profits excluding provisions) being relatively immaterial. This is consistent with findings in the literature that document the stability of bank net interest income (e.g., Drechsler et al. (2021)). Therefore, my mechanism appears to be the key driving force behind the fall in profits. The second step I would need to document is that the fall in profits (bottom-left panel of Figure 1.2) is *causing* most, if not all, of the increase in leverage (bottom-right panel of Figure 1.2). The accounting identity is less straightforward in this case as we are utilising information from both the income statement and balance sheet. The approach I take is to utilise the following identity for a balance sheet item at time t:

$$\frac{\text{Cumulative Profits}_t}{\text{Assets}_t} - \frac{\text{Cumulative Dividends}_t}{\text{Assets}_t} \approx \frac{\text{Equity}_t}{\text{Assets}_t} = \frac{1}{\text{Leverage}_t} \tag{1.3}$$

Note that (4.35) shows we need a measure of cumulative profits to transform an income statement measure (a flow) to a balance sheet measure (a stock). Equity at time t is constructed by adding all profits earned before t to the starting equity then subtracting all dividends paid before t and finally making some accounting adjustments (e.g., revaluations) at horizon t. While I do not have a direct measure of the accounting adjustments, I can construct the two cumulative measures: cumulative profits and cumulative dividends (accumulated from 1984 to $1984+t$). Note that the final term, equity divided by assets, is simply the inverse of leverage. Therefore, if my proposed mechanism is important, it should be the case that the variation in leverage (or the inverse of leverage) is driven by the variation in profits. Figure 1.4 below shows the impulse response of the first, second, and final term of (4.35) which are obtained by estimating (1.1) with each of those terms as the outcome variable z_t.

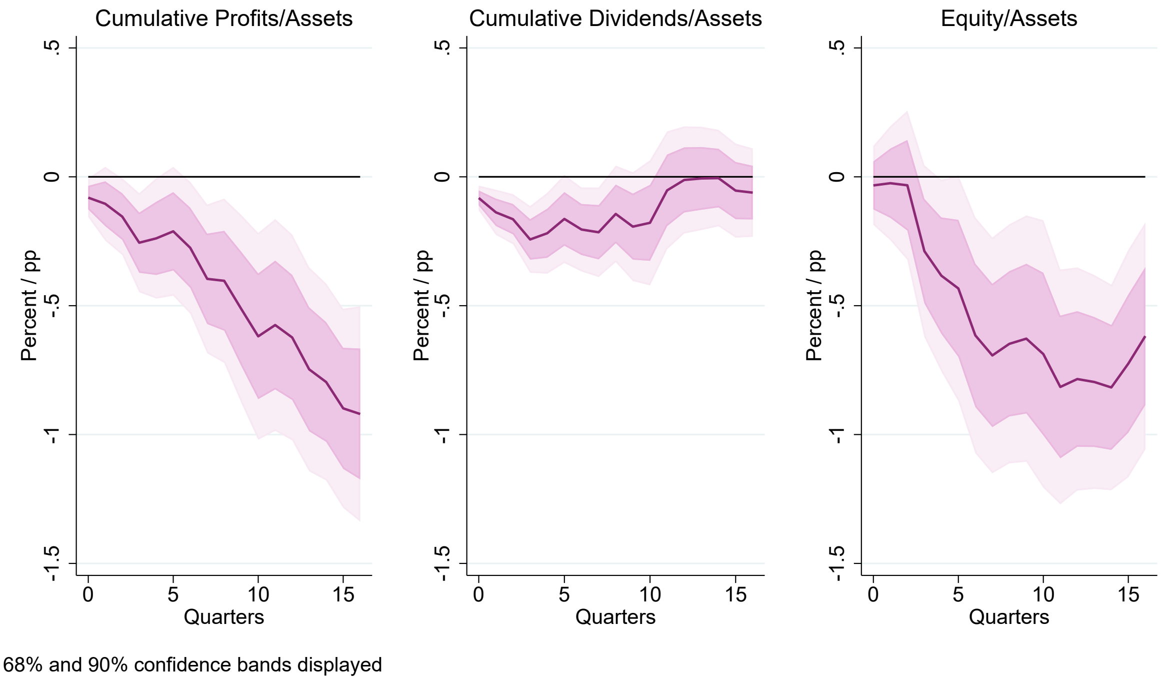

Figure 1.4. Decomposing the Leverage Increase

As can be seen, the variation in overall leverage (or more precisely the inverse of leverage) is largely driven by the variation in cumulative profits. As one might expect due to the potential penalties associated with reducing dividend payments (Guttman et al. (2010)), the response of cumulative dividends is muted. Moreover, while not shown, there is little unexplained variation after accounting for cumulative profits and dividends which implies that accounting adjustments would not be driving the overall response. Therefore, I have shown that my mechanism is driving the overall response in leverage as the decrease in profits is largely driven by the increase in loan losses and the increase in leverage is largely driven by the decrease in profits.

One take-away from this section thus far is that the macro-banking models used to understand monetary policy and its interaction with financial stability should allow for, at the very least, the potential for contractionary interest rates to raise bank leverage, given the robustness of the empirical moment. Furthermore, understanding what specifically drives loan losses will be key to determining the features that might be important when developing such models.

1.3.4 Drivers of Loan Losses

In Section 1.3.3, I showed that the rise in loan losses drives the variation in bank leverage in response to a contractionary monetary policy shock. However, it is not clear precisely why such a shock causes loan losses to rise in the first place. In this section, I attempt to shed light on this question using aggregate data.

Intuitively, one can think of contractionary monetary policy as leading to unexpected loan losses for two broad reasons. First, a higher FFR may directly raise the loan-servicing cost on floating-rate loans of any maturity or fixed-rate loans with a short maturity. This would reduce a borrower's ability to repay and hence raise loan losses. Second, a higher FFR may subsequently reduce incomes due to its recessionary impact as measured by a fall in GDP. This would also reduce a borrower's ability to repay and hence raise loan losses. However, a contractionary monetary policy shock both increases loan-servicing costs by directly raising the FFR and reduces borrower income by reducing GDP. Therefore, it is unclear by looking at such shocks whether a higher FFR or lower GDP is driving loan losses.

One approach to determine whether a higher FFR or lower GDP is driving loan losses is to consider variation that only affects one of the two factors. Cost-push shocks provide such variation as central banks are less likely to react to such shocks by raising interest rates. Therefore, cost-push shocks often feature little to no change in the FFR and hence minimal direct impact on loan-servicing costs but still have a decline in GDP and hence a reduction in borrower income. As highlighted in Section 1.2.4, oil shocks are a clear example of cost-push variation. Indeed, in response to an oil shock, we would expect a fall in GDP with little reaction of the FFR. This leads to the following empirical test: if loan losses are driven by the direct impact of the FFR on loan-servicing costs, then we expect loan-loss provisions to rise in response to a contractionary monetary shock but not in response to an oil shock.

Figure 1.5 presents the results of the empirical test. Specifically, Figure 1.5a shows the effect of a contractionary monetary policy shock on the FFR, GDP, loan-loss provisions and bank leverage while Figure 1.5b shows the effect of an oil shock on these same variables. As expected, both result in a decline in GDP, while only the monetary policy shock features a meaningful rise in the FFR. Interestingly, loan-loss provisions rise in response to the contractionary monetary policy shock but not in response to the oil shock which suggests that loan losses are more likely driven by the direct impact of the FFR on loan-servicing costs. Moreover, the fact that leverage only rises in the case of the monetary shock provides further validation of the importance of the loan-loss mechanism in driving variation in leverage.

Figure 1.5 also extends the empirical test with the risk perception series mentioned in Section 1.2.4. While the risk perception series is not a shock series, one can nonetheless use it as an additional robustness check. A rise in PSV (the measure of risk perception) can be thought of as a decline in risk and an increase in optimism by agents in the economy. Therefore, it behaves in a way similar to a positive aggregate demand shock. Figure 1.5c shows both the FFR and GDP rise, as is typically the case in response to a positive demand shock. Strikingly, loan-loss provisions rise despite the improvement in economic activity which suggests that the direct impact on loan-servicing costs due to the rise in the FFR is especially important.

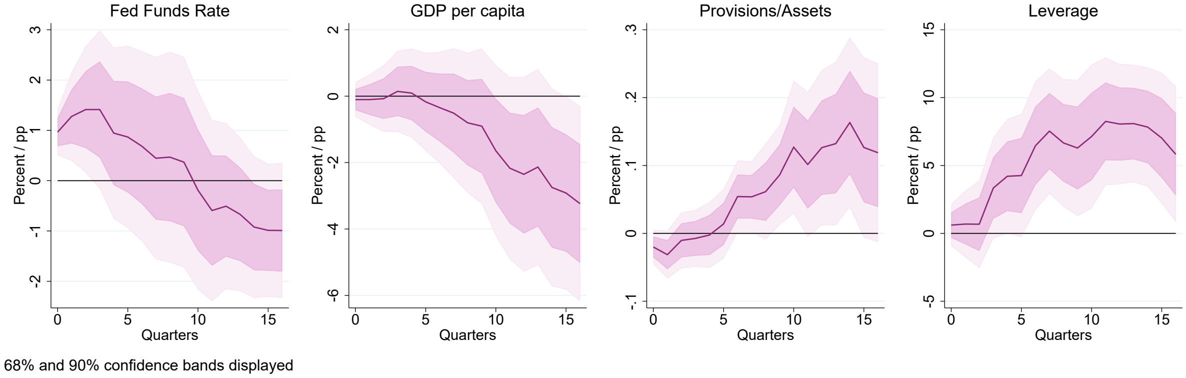

(a) Monetary Policy Shock

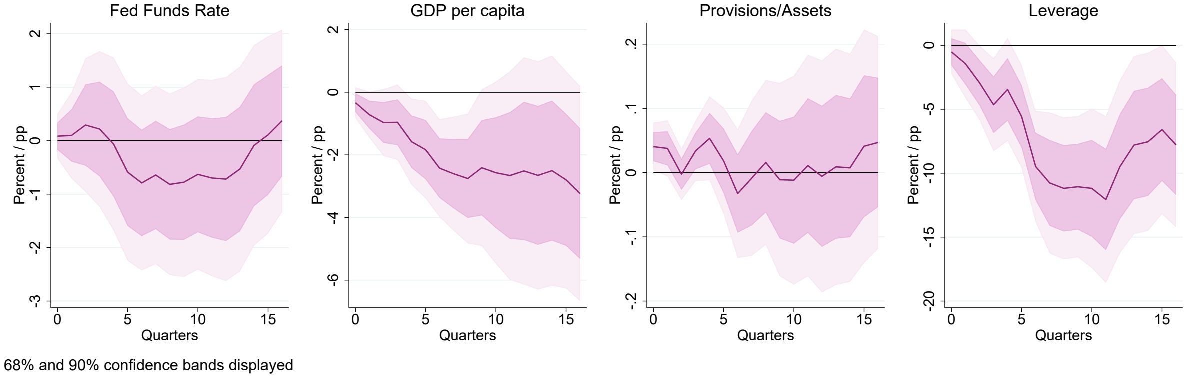

(b) Oil Shock

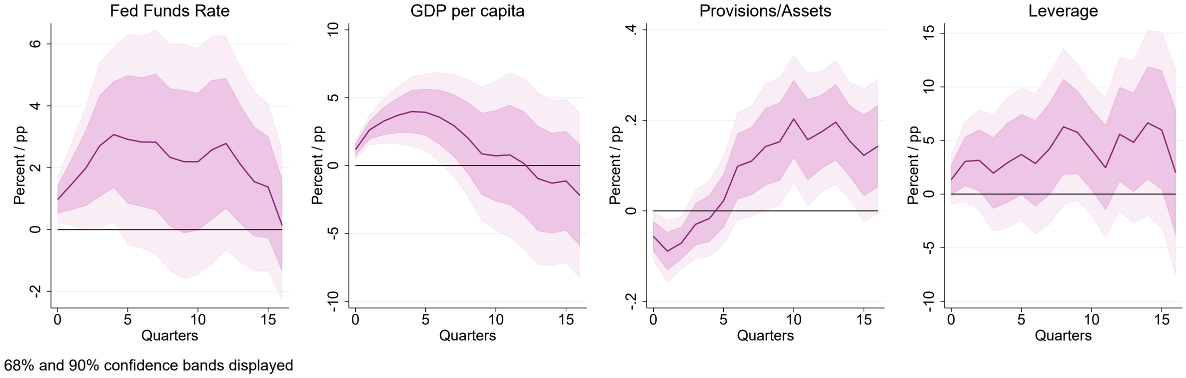

(c) Risk Perception (PVS)

Figure 1.5. Disentangling the Drivers of Loan Losses

Overall, Figure 1.5 lends support to the idea that loan losses are driven by higher interest rates directly increasing loan-servicing costs. Therefore, the overarching loan-loss mechanism is as follows: higher interest rates result in higher loan-servicing costs on loans more directly exposed to interest rates such as floating-rate loans. Loan losses appear to be driven by these higher loan-servicing costs. The increase in loan losses drives a decline in bank profits which reduces bank equity, and ultimately increases bank leverage. While the analysis highlights the potential role of floating-rate loans, it is still suggestive evidence. Therefore, in Section 1.5, I formalise the role of floating-rate loans in a simple banking model and in Section 1.6, I test it using microdata. Nonetheless, given I have now documented a clear mechanism, it is natural at this point to ask where the theory deviates from the empirical evidence whether this mechanism is captured in existing models.

1.4 Why do so many models generate a counterfactual leverage response?

In Section 1.3, I documented a robust finding: contractionary monetary policy shocks increase bank leverage. This result is almost entirely driven by the loan-loss mechanism whereby an unexpected increase in interest rates drives up loan losses at banks which reduces bank profits, subsequently eroding their equity, and ultimately increasing their leverage. I also provided suggestive evidence that floating-rate loans may play an important role in this mechanism. However, this mechanism, and the role of floating-rate loans in generating credit risk, is largely missing from the theoretical literature. In addition to missing the empirically dominant mechanism, much of the theoretical literature makes the opposite claim that leverage falls in response to a contractionary shock. While some do make an empirically consistent claim, they entirely ignore the loan-loss mechanism, and as a result have other predictions that are inconsistent with the observed data.

The divergence of the literature from the empirical evidence appears to derive from three broad, though not necessarily mutually exclusive, modelling choices: relying on a substitution effect; relying on procyclical leverage; and, relying on an a profitability channel. In this section, I will explain how each of these modelling choices leads the model to generate empirically inconsistent predictions as well as highlighting the types of papers in each category.

1.4.1 Models that rely on a substitution effect

The substitution effect is perhaps the most intuitive and simple mechanism that generates a counterfactual response of leverage to contractionary monetary policy shocks. Specifically, the substitution effect implies that an increase in the interest rate raises the relative cost of debt financing for banks and so banks substitute away from debt financing. A reduction in the reliance on debt financing is equivalent to a reduction in leverage. All else equal, this implies that higher rates reduce bank leverage, a claim I have shown to be empirically inaccurate.

Given the relative simplicity of the substitution effect, I will not go into the details of any particular model; rather I will briefly highlight some examples. The overarching message of these models is summarised in the review paper by Ajello et al. (2022): "Accommodative monetary policy reduces the cost of funding for banks, and thus may increase reliance on debt by banks."

This type of mechanism is common across the literature. For example, Angeloni and Faia (2013) introduce banks to a conventional DSGE model with nominal rigidities. Banks exist in the model because they can extract more liquidation value from projects. Banks are financed with deposits and equity and they are also subject to the risk of a run. The return on a project is equal to the expected value plus a random shock. Moreover, a run occurs if the outcome of a project is too low to repay depositors. If there is a contractionary monetary policy shock, the deposit rate increases which reduces the bank's ability to repay its depositors. This increases the probability of a run and so the bank reduces its deposits which decreases its leverage. Indeed, this mechanism is essentially a substitution effect that is induced by an endogenous run probability.

Another, albeit very different, example is the model by Drechsler et al. (2018b). They develop a dynamic asset pricing model in which monetary policy affects the risk premium component of the cost of capital. Risk-tolerant agents (banks) borrow from risk-averse agents by taking deposits to fund levered investments. Leverage exposes banks to funding shocks. As such, banks hold liquidity buffers composed of safe assets (e.g., US Treasuries) to insure against such funding shocks. If the central bank raises interest rates, it raises the liquidity premium because the cost of holding liquid securities increase. This increase in the price of funding shock insurance means banks will reduce their liquidity buffers. Therefore, with lower insurance, banks reduce their exposure to funding shocks by reducing deposits. Again, this is essentially a substitution effect but in this model it is induced by the dynamics of liquidity insurance.

For the substitution effect to increase leverage in response to a contractionary monetary policy shock, the following must be true: (i) debt liabilities fall; and (ii) the fall in debt liabilities is greater than the fall in equity. In Figure 1.6 below, I show that debt liabilities do fall in my data, consistent with (i). However, the more important contribution of my empirical analysis is that (ii) does not hold in the data. As I show in Section 1.3, the empirically dominant mechanism is the loan-loss mechanism which not only offsets the effect on leverage from falling debt liabilities, but actually leads to a reversal in sign such that contractionary monetary policy shocks increase leverage.

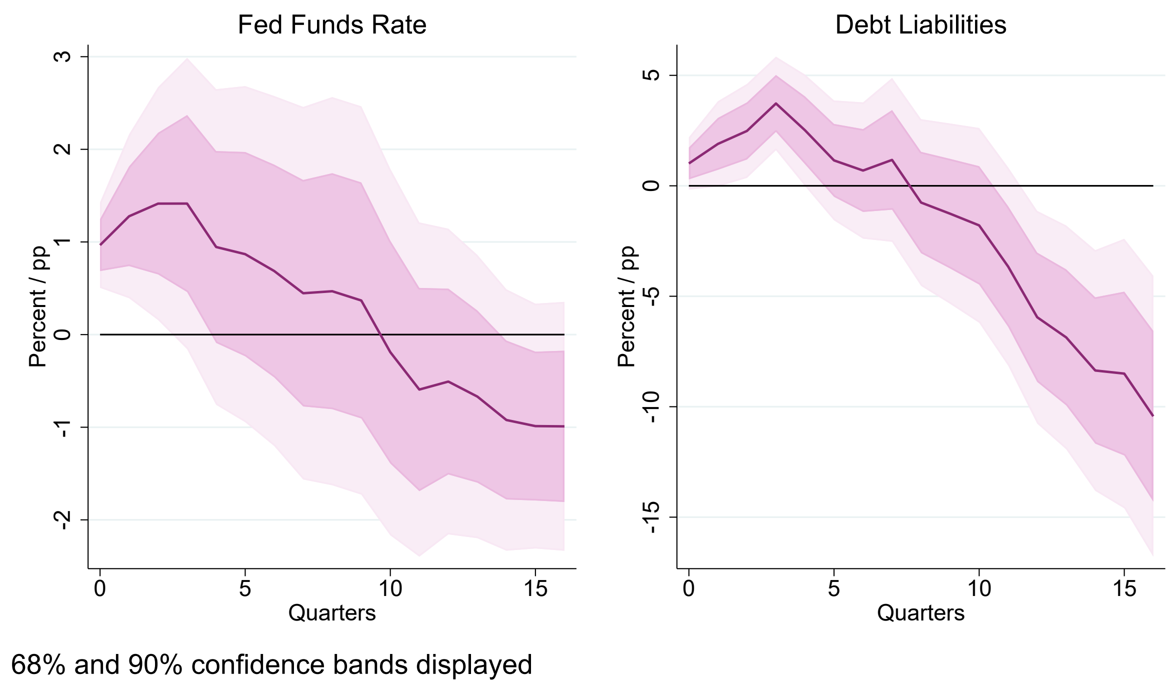

Figure 1.6. Impulse Response of Debt Liabilities to Contractionary Monetary Shock

1.4.2 Models that rely on leverage procyclicality

Like models based on the substitution effect, this class of models are similarly eclectic in their underlying structures, but have the common feature that the results are driven by the procyclicality of bank leverage. This procyclicality has been widely documented in the literature (e.g., Adrian and Shin (2010), Laux and Rauter (2017), and Adrian et al. (2019)). Such studies typically document this procyclicality by showing a positive relation between the growth of bank leverage and the growth of bank assets. The latter is considered procyclical as bank lending grows during a boom and shrinks during a bust. Figure 1.7 shows that leverage is indeed procyclical in my data. Specifically, it shows the positive correlation between the growth in bank leverage and the growth in bank assets as well as growth in bank leverage and GDP growth directly.[17]

[17] One can also do a simple regression of the growth in leverage on GDP growth which would yield a positive coefficient with a t-statistic of 1.87.

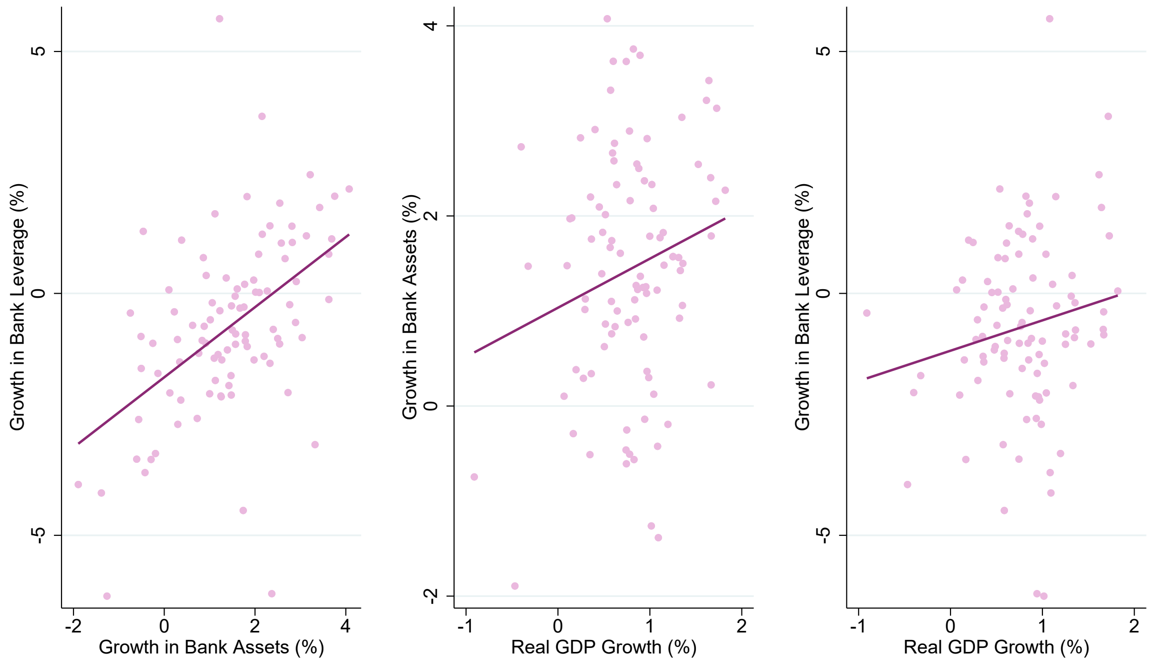

Figure 1.7. Procyclicality of Bank Leverage

In these models, a contractionary monetary shock will reduce output and because leverage is procyclical, it will also reduce leverage. Such models rely on leverage procyclicality in different ways. Some use procyclicality of leverage as a target or a measure of success of the model. For example, Rannenberg (2016) points out that by introducing a firm sector in the spirit of Bernanke et al. (1999) to the model of Gertler and Karadi (2011) (he terms this combined model the "full model"), he is able to generate procyclical leverage. Specifically, he concludes that "in the full model, bank leverage declines in response to contractionary monetary policy and productivity shocks, which allows the full model to match the procyclicality of bank leverage in U.S. data. By contrast, bank leverage in the Gertler–Karadi-type model is strongly countercyclical."

However, one cannot match conditional moments in a model to unconditional moments in the data as these are two entirely different measures. Consider the evidence in Figure 1.5. Both negative oil shocks and contractionary monetary policy shocks reduce GDP. While the former decreases leverage, the latter increases it. One cannot conclude whether leverage is unconditionally procyclical or countercyclical from this information alone. Indeed, looking at the correlations of leverage with GDP from these impulse responses alone would result in the conclusion that leverage is both procyclical and countercyclical. Moreover, as Galí (1999) points out, evaluating models based on their ability to match unconditional moments in the data can be highly misleading as the model may perform well according to that criterion despite providing a very distorted image of the economy's response to different types of shocks. Therefore, a conditional leverage moment, as I have documented, serves as a much sharper test of the model, and one that directly provides insight on the role of monetary policy, though one must ensure that they make like-for-like comparisons.

Like Rannenberg (2016), many papers do not distinguish between the procyclicality of leverage in the data (an unconditional moment) and the response of leverage to monetary policy (a conditional moment). This leads to the conclusion that monetary policy should 'lean against the wind' by tightening in response to increasing leverage. One particularly prominent, albeit highly stylised, paper that makes this type of argument is Woodford (2012). He provides a simple and reduced-form model of the way in which endogenous state variables affect the probability of a crisis and what this means for optimal monetary policy. To highlight in more detail how issues arise when using this type of procyclicality, I will focus on the set-up in Woodford (2012). The advantage of this model is in its simplicity which allows one to easily see the intuition.

The model is a fairly typical three-equation New Keynesian (NK) model except with two types of households: those that are credit constrained and those that are not. This is represented by the existence of a credit friction Ω_t which essentially measures the gap at any point in time between the marginal utilities of the two types of households. Woodford (2012) then derives a modified intertemporal IS equation:

$$y_t - g_t + \chi\Omega_t = \mathbb{E}_t[y_{t+1} - g_{t+1} + \chi\Omega_{t+1}] - \sigma[i_t - \mathbb{E}_t\pi_{t+1}] \tag{1.4}$$

where y_t is the output gap, g_t is government purchases, i_t is the nominal interest rate set by the central bank, π_{t+1} measures inflation between period t and $t+1$, and the coefficients satisfy $\chi, \sigma > 0$. All variables represent deviations from the steady state. The only difference between equation (1.4) and the standard IS equation is the credit friction. Indeed, as one would expect, a higher credit friction would behave similarly to the effects of a reduction in government purchases. Therefore, real aggregate demand now also depends on the severity of credit frictions in the economy. A similar approach yields a modified NK Phillips curve.

$$\pi_t = \kappa_y y_t + \kappa_\Omega \Omega_t + \beta \mathbb{E}_t \pi_{t+1} + u_t \tag{1.5}$$

where u_t is a composite term denoting the different exogenous cost-push factors. Again, the Phillips curve is exactly the same as that in the standard model except for the additional credit friction. A key component of the model is to incorporate some endogeneity in how the credit fiction evolves. Ω_t is assumed to always be in one of two states: a normal state (low value of Ω_t) and a crisis state (high value of Ω_t). Each period, the probability of entering the normal state when in a crisis state is δ, while the probability of entering the crisis state when in a normal state is γ_t which is an increasing function of bank leverage (L_t). Intuitively, as leverage is higher, the probability of going into a crisis is higher. Therefore, to complete the model, Woodford (2012) connects leverage with the remaining endogenous variables by postulating a simple law of motion:

$$L_t = \rho L_{t-1} + \xi y_t + v_t \tag{1.6}$$

where v_t represents an exogenous disturbance term and importantly ξ is assumed to be positive. Therefore, this law of motion embeds the procyclicality of leverage as leverage is an increasing function of the output gap. Indeed, this type of assumption is the core reason models in this class are unable to generate empirically consistent dynamics.

To complete the framework, Woodford (2012) assumes that the goal of policy is to minimise the following loss function:

$$\frac{1}{2}E_0\sum_{t=0}^{\infty}\beta^t\left[\pi_t^2+\lambda_y y_t^2+\lambda_\Omega\Omega_t^2\right] \tag{1.7}$$

This is an intuitive form of the loss function as the central bank is simply minimising losses from inflation, output, and financial instability. However, the problem arises because of the way in which monetary policy and leverage now intertwine. A contractionary monetary shock will reduce the output gap as is typically the case. However, because of equation (1.6), the same shock will also reduce leverage. Indeed, Woodford (2012) concludes that the model implies one should use monetary policy to 'lean against' a credit boom (which in this model would be to reduce leverage) even if it requires missing target values for inflation and the output gap. In this model, the primary prediction is inconsistent with my empirical findings and as such the consequences of following such a rule are more severe. For example, consider a central bank that is following such a rule when inflation and the output gap are on target, but credit frictions are far too high. As Woodford (2012) mentions, it would be appropriate for the central bank to use contractionary monetary policy. The consequences would not only see both inflation and the output gap falling below target, but leverage would actually rise due to increasing loan losses which are missing from the model. This would unambiguously worsen losses according to the central bank loss function.

The Woodford (2012) model, while highly stylised, is very influential as it builds on the workhorse NK structure. However, it entirely misses the empirical loan-loss mechanism and instead relies on a postulated law of motion that embeds procyclicality. This procyclicality between leverage and output is correlational, not structural. Indeed, when models rely on procyclicality in this way, they typically argue that monetary policy should lean against the wind because contractionary shocks, contrary to the evidence in this paper, reduce bank leverage.

1.4.3 Models that rely on a profitability channel

In this class of models, profitability and leverage move together and are connected by an incentive compatibility leverage constraint. Furthermore, these models have the feature that *any* negative shock will increase bank profitability as well as bank leverage. As such, while these models correctly show that leverage increases in response to a contractionary monetary policy shock, they have leverage increasing alongside an increase in profitability (which I term the profitability channel). This is at odds with the empirical evidence that the increase in leverage following a contractionary monetary policy shock *is caused by* a decrease in profits.

This class of models build on the canonical models of Gertler and Kiyotaki (2010) and Gertler and Karadi (2011) which are some of the most influential macroeconomic models featuring a banking sector. The defining feature of this class is that they use a Gertler-Karadi-Kiyotaki-type constraint (i.e., an incentive-compatibility leverage constraint) to model banks which generates the profitability channel.[18]

Given the empirical inconsistencies arising from using a Gertler-Karadi-Kiyotaki-type constraint are intricate, I will focus on the specific set-up in Gertler and Karadi (2011) to highlight how the issues arise. I choose Gertler and Karadi (2011) for two reasons. First, given it is one of the foundational models, most models in this class typically have the same underlying structure. Second, while Gertler and Kiyotaki (2010) is also a foundational model, Gertler and Karadi (2011) incorporates nominal rigidities and so is better able to highlight the impact of monetary policy on bank leverage.[19]

Gertler and Karadi (2011) builds on the seminal monetary dynamic stochastic general equilibrium (DSGE) models of Christiano et al. (2005) and Smets and Wouters (2007) by incorporating banks that transfer funds between households and non-financial firms. Banks exist

[18] This modelling approach is widely used in the literature, e.g., Gertler and Kiyotaki (2015), Maggiori (2017), Gertler et al. (2020), Ghote (2021), and Sims and Wu (2021).

[19] While Gertler and Kiyotaki (2010) is a purely real model, both models have the same underlying structure; one can think of Gertler and Karadi (2011) as extending the Gertler and Kiyotaki (2010) model to allow for nominal rigidities.

as they have expertise in evaluating and monitoring borrowers and a simple agency problem between banks and households constrains the ability of banks to raise deposits. The model features five different agents: households, goods producers, capital producers, monopolistically competitive retailers, and banks. Monetary policy is characterised with a simple Taylor rule. Without banks, the model is isomorphic to Christiano et al. (2005) and Smets and Wouters (2007). While the model is a sophisticated general equilibrium (GE) model, one need only analyse the banking block of the model to understand how Gertler and Karadi (2011) generate the correct leverage prediction as well as where the model is diverges from the empirical evidence. Therefore, I will focus on the partial equilibrium of the banking block to highlight the intuition and precisely depict the underlying mechanisms. For completeness, I will also show the results from simulating the full GE model to show that the core insights obtained from examining the banking block do not change once we account for GE dynamics.

I will follow a stylised version of the Gertler and Karadi (2011) model. Banks obtain deposits, B, from households. These funds are then 'lent' to non-financial firms which gives banks a claim on those firms where S depicts the quantity of those claims.[20] Each claim has price Q. Therefore, the net worth (equity), N, of the bank is given by the following balance sheet constraint:

$$N = QS - B \tag{1.8}$$

The stochastic return on a single unit of lending is R_k while a single unit of deposits pay a non-contingent return R. Both returns are determined endogenously. Given this structure, the bank's objective is to maximise the expected value of the bank, V, which is simply maximising the difference between the expected earnings on assets and interest payments on liabilities. The

[20] Technically, these loans by banks to non-financial firms are perfectly state-contingent debt and so are better thought of as equity.

value of the bank is therefore given by the following:

$$V = R_k QS - RB \tag{1.9}$$

We can plug in the balance sheet constraint, equation (1.8), into the bank objective function above to yield the following:

$$\begin{aligned} V &= R_k QS - R(QS - N) \\ &= \underbrace{(R_k - R)}_{\text{profitability}} QS + RN \end{aligned} \tag{1.10}$$

Equation (1.10) shows that a bank's value is a function of the premium the bank earns on its assets, which I have termed profitability. We can already see that in this model there is no measure of loan losses that were key to the empirical mechanism documented in Section 1.3. While one could argue that loan losses might already be included in the endogenously determined R_k, I will explain how this is not the case.

Thus far, the model is fairly standard. However, an important feature of equation (1.10) is that so long as the bank has positive profitability (i.e., $R_k - R > 0$), it will want to infinitely expand its assets. Put differently, bank value V is increasing in assets when banks have positive profitability. Therefore, a core component of this class of models is the introduction of a moral hazard/costly enforcement problem which generates an endogenous leverage constraint (i.e., the Gertler-Karadi-Kiyotaki constraint) and thus prevents banks from infinite expansion. The costly enforcement problem is modelled as follows. After households place their deposits in a bank, the bank can divert a fraction λ of the deposits for itself. However, if the bank diverts those deposits, the depositors will force the bank into bankruptcy and recover the remaining $1 - \lambda$ share of assets. Therefore, rational depositors will only deposit at a bank if the bank has no incentive to divert assets. This yields the following incentive constraint which must be satisfied:

$$V \geq \lambda QS \tag{1.11}$$

Intuitively, the incentive constraint above is saying that a depositor would only deposit at a bank if the bank value (i.e., the value the bank obtains from being honest) is greater than the value the bank receives if it diverts assets (i.e., the value the bank obtains from not being honest). One can already see that banks with high value will be able to attract more deposits and subsequently grow their assets. Therefore, the incentive constraint prevents banks from expanding their assets infinitely as they need their value, V, to be larger than the share of divertible assets. As such, banks will expand up to that point (so long as profitability is positive). This implies that equation (1.11) will hold with equality and so we can equate equations (1.10) and (1.11) together.[21] This yields the following:

$$\begin{aligned} \lambda QS &= (R_k - R)QS + RN \\ \implies \text{leverage} \equiv \frac{QS}{N} &= \frac{R}{\lambda - (R_k - R)} \end{aligned} \tag{1.12}$$

Now we have an equation for leverage (i.e., total assets divided by net worth). This equation, which is derived from the bank problem and incentive constraint alone, has a very important implication: leverage is increasing in profitability (where profitability is $R_k - R$). The intuition behind this implication is that if a bank is able to make more profits, then it has less incentive to divert assets and cheat depositors. As such, depositors are more willing to lend to the bank which enables the bank to increase its leverage. Note that because of equity issuance frictions, the adjustments come from leverage. We are now in a position to contrast this simple intuition to the empirical findings.

Recall, I show that a contractionary monetary policy shock reduces profits which depletes

[21] Gertler and Karadi (2011) explicitly state that the constraint always binds within a local region of the steady state.

net worth and subsequently increases leverage. Note specifically that leverage is rising *because* of falling profits. On the other hand, the model has leverage rise together with profits rising (i.e., the profitability channel). This is an important inconsistency. The reason profits fall in the data is driven primarily from a rise in loan losses as borrowers with floating-rate loans are less able to repay. The model has no measure of loan losses and as such does not capture that profits fall following a contractionary monetary policy shock. Therefore, even though the interest rate on lending, R_k, is endogenously determined, it is unable to capture the loan-loss dynamic. Hence, modelling banks through this type of Gertler-Karadi-Kiyotaki incentive constraint generates an empirically inconsistent profitability channel. While my empirical results were based on the monetary policy shocks identified in Romer and Romer (2004), one can actually use the monetary policy shock series in Gertler and Karadi (2015) to see whether the Gertler and Karadi (2011) model would be consistent with the Gertler and Karadi (2015) shock series.

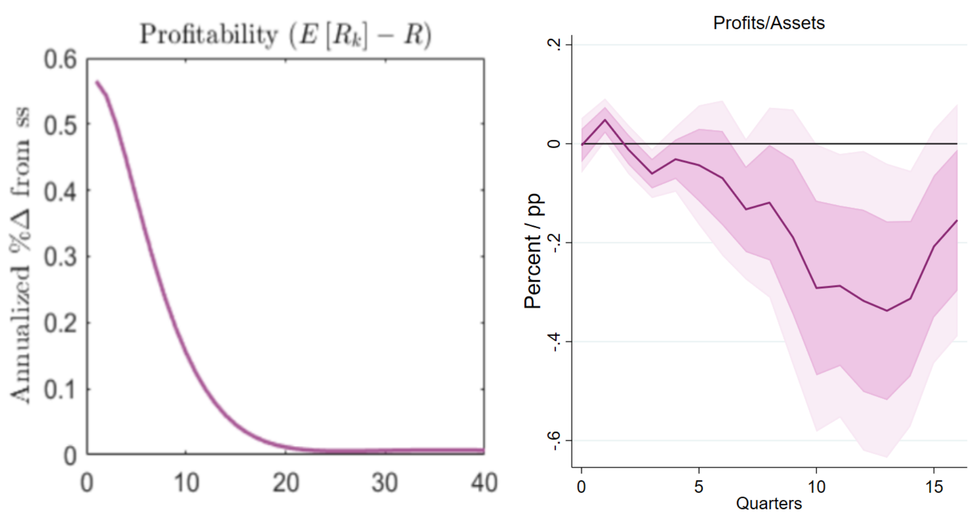

(a) Model: Gertler and Karadi (2011)　　**(b)** Empirics: Gertler and Karadi (2015)

Figure 1.8. Impulse Response of Profits to Contractionary Monetary Shock

Figure 1.8a shows that in the Gertler and Karadi (2011) model, profits increase following a contractionary monetary policy shock, before returning to steady state. However, Figure 1.8b, which uses the same specification as throughout this paper, shows the shock series in Gertler and

Karadi (2015) predicts a decrease in profits following a contractionary shock.[22] One possible explanation for the inconsistency with respect to profits is that $R_k - R$ can represent several different measures of profitability, where the book profits of banks is one possible measure. Other measures could include the risk premium in the economy or the expected return on bank stocks.

Moreover, one could argue that Gertler and Karadi (2011) and similar models are still able to predict that leverage rises following a contractionary monetary policy shock and perhaps that is sufficient despite the mechanism being empirically inconsistent. There are two important problems with this line of reasoning. First and foremost, one important rationale for developing macroeconomic models with a microfounded banking sector such as Gertler and Karadi (2011) is to help us understand the underlying economic mechanism. However, not only does the model miss an important channel through which monetary policy is being transmitted to banks (i.e., loan losses), it suggests a mechanism that is contradicted by the data as profits rise in the theoretical model of Gertler and Karadi (2011) but fall according to the Gertler and Karadi (2015) empirical evidence.

Secondly, even if the primary objective of this class of models is prediction rather than capturing the underlying economic mechanism, the results may still be misleading. To highlight why, let us return to the Gertler and Karadi (2011) model and rewrite the definition of leverage using the incentive constraint (i.e., using equation (1.11) to replace QS). This yields the following:

$$\text{leverage} = \frac{V}{\lambda N} \tag{1.13}$$

[22] Given that the GK shock starts later than the RR shock, the data underlying the figure is from 1994 onwards and eight lags are used instead of sixteen due the shorter time horizon. The empirical measure of profits is profit divided by assets and so is measuring profitability in a way that closely resembles the measure of profitability in the model.

In these types of models, bank value (V) is linear in bank assets and net worth[23]

$$V = vQS + \eta N \tag{1.14}$$

where v is the marginal value of expanding assets and η is the marginal value of expanding net worth. Plugging equation (1.14) into equation (1.13) yields the following:

$$\begin{aligned} \text{leverage} &= \frac{vQS + \eta N}{\lambda N} \\ \implies \lambda \cdot \text{leverage} &= v\frac{QS}{N} + \eta \\ \therefore \text{leverage} &= \frac{\eta}{\lambda - v} \end{aligned} \tag{1.15}$$

where the last step makes use of the fact that leverage $\equiv \frac{QS}{N}$. Equation (1.15) highlights a very important implication of this type model structure: leverage is increasing in the marginal value of net worth (as well as the marginal value of assets). Given the equity frictions, *any* negative shock that increases the marginal value of net worth (η) and assets (ν) will also increase leverage.[24] Therefore, while this feature allows such models to correctly predict that leverage rises in response to a contractionary monetary policy shock, they also predict that all other negative shocks would yield a rise in leverage which is a much stronger claim. If this were true, then perhaps one could put less emphasis on the underlying economic mechanism. However, I have already provided one counterexample to the claim that any negative shock will increase bank leverage. Figure 1.5b shows a negative oil shock that leads to a decrease in leverage. So why are Gertler and Karadi (2011) able to accurately predict leverage in the case of the contractionary monetary policy shock scenario but not negative oil shock? One reason is because there is no

[23] Indeed, as highlighted in Ghote (2021), having V proportional to net worth implies that the bank problem is scale invariant. As such, optimality implies that leverage is the same for all banks regardless of their net worth which allows for a representative bank.

[24] A negative shock increases the marginal value of net worth because it causes an on-impact decrease in the price of capital, Q. This reduces bank net worth as bank assets are now worth less. However, a decline in net worth means banks are less able to lend which decreases total loans. A decline in total lending raises the expected profitability of lending which raises the marginal value of net worth.

distinction by type of shock in the model. However, as is evident in Figure 1.5, the type of shock matters empirically not just in terms of magnitude but also direction. While both an oil shock and monetary policy shock cause a decline in GDP, only the latter increases loan losses, which also highlights the role of floating-rate loans. This again underscores the importance of the underlying loan-loss mechanism; a component missing from Gertler and Karadi (2011). Ensuring the mechanism is modelled appropriately ensures that the predictions are made in the right context. In this sense, the prediction in this class of models may be misleading.

Thus far, we have only used a partial equilibrium analysis to understand how and why models that use a Gertler-Karadi-Kiyotaki type constraint to model banks generate implications that are inconsistent with empirical evidence. While such analysis more easily highlights the underlying intuition and dynamics, one might argue that the full GE model could generate different results. Therefore, in Figure 1.9 below, I show the results obtained from the full GE model in response to a contractionary monetary policy shock.[25] As can be seen, in the full GE model, a contractionary shock decreases output, increases profitability, increases the marginal value of net worth, and increases leverage, all claims that were evident from analysing the banking block alone. Moreover, consistent with equation (1.15), Gertler and Karadi (2011) also show that a negative total factor productivity shock leads to an increase in leverage.

[25] The model code is obtained from the Macroeconomic Model Data Base (see https://www.macromodelbase.com/).

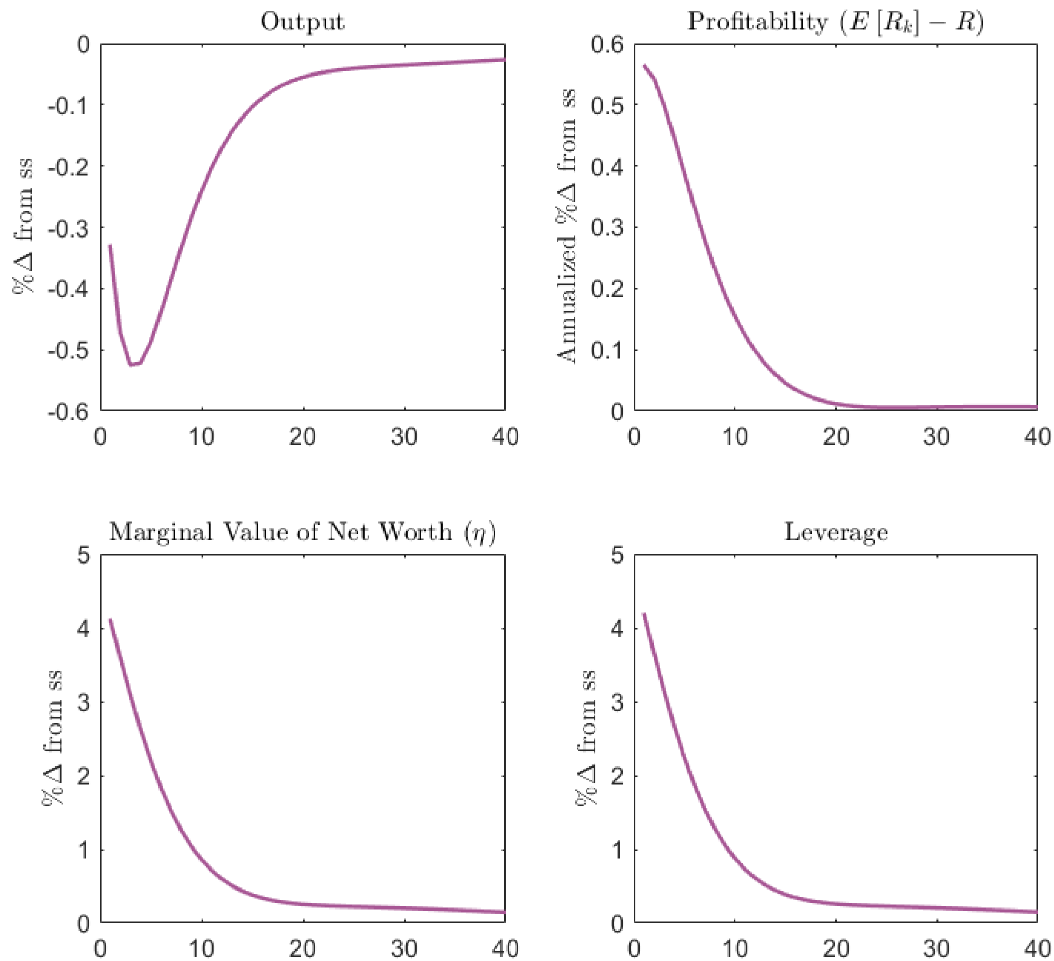

Figure 1.9. Gertler and Karadi (2011) Model Response to Monetary Policy Contraction

Therefore, for understanding bank leverage, the set of results in this section highlights that the insights obtained from examining the structure of the banking sector alone survives GE dynamics. Indeed, given the implications of my analysis of the banking block, and that the remaining components of these model (i.e., households and firms) are fairly standard, it appears that the inconsistencies in such models arise primarily due to the profitability channel, which ignores the loan-loss mechanism.

One important aside here is that there are typically two types of analyses. One that considers exogenous monetary policy shocks (e.g., Gertler and Karadi (2011) and Drechsler et al.

(2018b)) and another that derives a monetary policy *rule* that suggests leaning against the wind (e.g. Woodford (2012) and Ghote (2021)). Therefore, one may argue that the analysis in this paper is only pertinent to the papers considering exogenous shocks. However, Wolf and McKay (2023) show that analyses using exogenous shocks and analyses considering alternate policy rules can be equivalent under certain conditions. Specifically, if policy affects private-sector behaviour only through the current and future expected path of the policy instrument (the case in most models), then in the eyes of the private sector, a prevailing non-leaning-against-the-wind monetary policy rule subject to a particular sequence of contractionary interest rate shocks is identical to some counterfactual leaning-against-the-wind policy rule. Put simply, the private sector is not able to distinguish between a contractionary shock and a change in the monetary policy rule that would generate the same contractionary shock. As such, my empirical findings are relevant for analyses involving both exogenous shocks and modified policy rules, albeit more directly for the former.[26]

1.5 An Empirically Consistent Theoretical Model

The models presented thus far capture a wide variety of the results in the literature as they are some of the most foundational models. However, they all fail to capture the empirical dynamics that I have documented. A crucial missing ingredient is the loan-loss mechanism. Interestingly, while many of the models explored in the previous section are GE models (as is typically the case when modelling monetary policy), we did not need to examine the whole GE structure to see where problems arise. Indeed, the problems shown in Section 1.4 arise primarily from how the banking system is modelled (e.g., from the bank problem in Gertler and Karadi (2011) or the bank leverage law of motion in Woodford (2012)).

It is instructive to consider whether partial equilibrium banking models, in particular, those that do not rely primarily on a substitution effect or leverage procyclicality, can be more

[26] Wolf and McKay (2023) note that their result is less suited to study policies that alter the steady state (e.g., changes in the inflation target). However, many analyses of optimal rules compare different cyclical stabilization policies such as augmented Taylor rules, where the results of Wolf and McKay (2023) apply.

empirically consistent. A few examples that are able to generate mostly consistent empirical dynamics are Van den Heuvel (2009) and Corbae and Levine (2022). Van den Heuvel (2009) develops the bank capital channel of monetary policy which sees profits fall and leverage rise following contractionary monetary policy. The underlying mechanism of the model is through maturity transformation rather than loan losses. However, he also shows how a default shock works in the model and it generates dynamics similar to the loan-loss mechanism. While the default shocks are exogenous, if they were a function of a contractionary monetary policy shock, the dynamics would appear to match those in the loan-loss mechanism. Nonetheless, there is still no explicit role for floating-rate loans which appear to be an important feature in the data (see Section 1.6). Corbae and Levine (2022) take a different modelling approach but also see profits fall and leverage rise in response to contractionary monetary policy. The mechanism is also different to the loan-loss mechanism as higher rates raise the marginal cost of financing for banks which induces greater risk-taking and a fall in profits. While such models get fairly close to matching the empirical dynamics, the mechanism underpinning the fall in profits in these models is different to loan losses and does not feature a role for floating-rate loans, both of which appear important for the empirical mechanism.

Therefore, from examining theories that are unable to match the empirical dynamics as well as those that match them better, one can surmise the following. First, the general equilibrium structure does not appear to be especially important in generating the dynamics of leverage in response to contractionary shocks. This means one can focus on a partial equilibrium banking model in order to illuminate the mechanism more clearly. Second, a common missing ingredient across most models is the loan-loss mechanism. As such, the model needs to capture loan losses that are increasing in contractionary monetary policy shocks. Third, there needs to be a potential role for floating-rate loans in generating the loan-loss mechanism. For these reasons, I develop a model with these three components. The purpose of the model is twofold. First, it formalises the role of loan losses in determining the response of bank profitability to contractionary monetary policy shocks and sheds light on the role of floating-rate loans. Second, it generates implications which I can then explore using microdata.

My model provides a novel way for thinking about banks by emphasising what I call *risk transformation* which works as follows. Banks are exposed to interest rate risk because their deposits are floating-rate liabilities (i.e., when interest rates rise, deposits become more expensive). To hedge the interest rate risk and alleviate the cash flow mismatch on their balance sheets, banks issue floating-rate loans. So when interest rates unexpectedly rise, while banks have to pay more to depositors, they also receive more income from floating-rate borrowers. However, this hedging strategy works by transferring the risk from banks to borrowers. Unlike banks, borrowers cannot hedge against unexpected interest rate changes.[27] As such, in response to contractionary monetary policy shocks, borrowers are less able to repay their loans which leads to loan losses for banks. Such losses represent a credit risk for the bank. Therefore, through issuing floating-rate loans, banks are conducting risk transformation as they are hedging interest rate risk at the expense of greater credit risk.

My model builds on Kirti (2020) but incorporates credit risk via loan losses. Consider a one-period model with the following timeline. First, banks make loans funded by deposits and internal net worth. Second, the realization of the monetary shock takes place. Finally, repayment occurs. Banks are exogenously endowed with deposits D and a loan portfolio of size L, as such, internal net worth is $N = L - D$. The key choice for banks is the share of floating-rate loans f_L in their loan portfolio.[28] The deposits are floating-rate liabilities. However, as shown by Drechsler et al. (2017), there is not perfect pass-through of the central bank interest rate to deposit rates. In the model, the pass-through coefficient, known as the deposit-beta (β^{dep}), is exogenous but one can microfound this by using the approach in Drechsler et al. (2017). The interest rate is a random variable $r = \bar{r} + \varepsilon$ where $\varepsilon \sim \mathcal{N}(0, \sigma^2)$. Therefore, $E[r] = \bar{r}, Var[r] = \sigma^2$. Note that ε

[27] There is an important distinction between unexpected and expected changes in interest rates. Because expected interest rate changes are procyclical, borrowers are naturally hedged as while their loan-servicing costs rise, they also receive greater cash flows. See Figure 1.15 and the associated discussion for further detail.

[28] One can think of this as a two-stage optimisation problem for banks. In the first stage, they choose the optimal balance sheet size (loans and deposits). In the second stage, they choose the share of their loans to be fixed or floating. Given I am only interested in the second stage, the loan size and deposits are exogenous.

is the monetary policy shock.

Note that the one-period nature of the model implies leverage will move in lockstep with profits. Specifically, if profits fall, leverage will rise. I choose this approach because, as in most models, banks have limited scope to adjust dividends or raise equity. Therefore, the interesting variation comes from the response of profitability which subsequently determines leverage. In my model, all of the variation in leverage in response to contractionary monetary policy shocks will be explained by profitability which is broadly consistent with my empirical findings.

I model banks as having risk-averse preferences in order for them to dislike risk.[29] As such, banks maximise value V by choosing the share of its loan portfolio that is floating rate:

$$\max_{f_L} V = g(\pi)$$

where π is bank profits and $g(\cdot)$ represents a risk-averse functional form. For analytical simplicity, I will use mean-variance preferences. So the bank objective function is

$$\max_{f_L} V = E[\pi] - \frac{\gamma}{2} Var[\pi]$$

where γ captures risk-aversion. Solving this yields an expression for f_L^* in terms of μ which is solved in equilibrium with the firm problem (see Appendix A.4 for the full model and an analytical expression for f_L^*). However, the core insight from the model comes from the following thought experiment: given the optimal choice f_L^*, what is a bank's profits and expected profits?

$$\pi = \overbrace{\underbrace{L(1-f_L^*)(\bar{r}+\mu^*(f_L^*))}_{\text{fixed-rate income}} + \underbrace{Lf_L^*(\bar{r}+\varepsilon+\mu^*(f_L^*))}_{\text{floating-rate income}} - \underbrace{D(\bar{r}+\beta^{dep}\varepsilon)}_{\text{cost of deposits}}}^{\text{Net Interest Income}} - \overbrace{Lf_L^*\theta(\varepsilon)}^{\text{Loan Losses}} \tag{1.16}$$

$$E[\pi] = L(1-f_L^*)\bar{r} + Lf_L^*\bar{r} + L\mu^*(f_L^*) - D\bar{r} - Lf_L^*E[\theta(\varepsilon)] \tag{1.17}$$

[29] The assumption that banks have risk-averse preferences is not uncommon in the literature. See for example Di Tella and Kurlat (2021).

μ^* is the equilibrium loan spread between the lending rate charged to firms and the central bank interest rate. Note that μ^* will be decreasing in the floating share of loans as banks have to accept a lower spread because firms are risk-averse and will also want to avoid bearing interest rate risk. One key point in (1.16) is that net interest income is not the same as profits. Indeed, as seen in the aggregate data, loan losses drive the vast majority of the variation in bank profits in response to a monetary policy shock. Any expected losses would already be priced in and therefore, to highlight the mechanism, I focus on the unexpected losses. The model only features unexpected loan losses on floating-rate loans. $\theta(\varepsilon)$ is the loan-loss rate where $\theta'(\varepsilon) > 0$ and $\theta'(\varepsilon)$ is linear in ε. This is intended to capture in a reduced-form way that loan losses are increasing in the size of the monetary policy shock. There are no unexpected loan losses on fixed-rate loans as a change in the central bank interest rate does not impact the loan-servicing cost of the fixed-rate borrower. This argument also rules out a recessionary channel of defaults as the purpose is to specifically highlight the role of floating-rate loans in order to explain that in the aggregate data we see loan losses rise with contractionary monetary policy shocks but not for other contractionary shocks. Note that this does not imply that recessions cannot cause loan losses but it simply highlights the loan-loss mechanism that appears to be induced by floating-rate loans.

Using (1.16) and (1.17), I define deviations from expected profitability (as measured by return on assets) as the following:

$$\begin{aligned} \Delta &= \frac{\pi}{L} - \frac{E[\pi]}{L} \\ \implies \Delta &= \underbrace{f_L^* \varepsilon - \frac{D}{L}\beta^{dep}\varepsilon}_{\text{interest rate risk}} - \underbrace{f_L^* \left(\theta(\varepsilon) - E[\theta(\varepsilon)]\right)}_{\text{credit risk}} \end{aligned} \tag{1.18}$$

Equation (1.18) represents a key insight of the model. The bank is exposed to interest rate risk because a contractionary monetary policy shock makes deposits more expensive. Floating-rate loans generate more revenue for the bank when interest rates increase and therefore banks issue floating-rate loans as a way to hedge interest rate risk. This is consistent with Kirti (2020) who shows empirically that banks that have a higher deposit pass-through (higher β^{dep}) issue more floating-rate loans. However, the core insight of (1.18) is that this interest rate risk hedge comes at the expense of credit risk. Specifically, the bank hedges the interest rate risk by passing that risk onto the borrower. If a borrower cannot hedge this risk, this generates loan losses for the bank. In the model, this is captured by $\theta(\varepsilon)$. A simple example can illustrate this more clearly. Consider a bank that issues a floating-rate loan that exactly tracks the central bank rate. If the central bank raises the interest rate, the borrower now has to pay more on the loan which raises the probability of default of the borrower. The bank has merely traded interest rate risk for credit risk. While in many models, banks do maturity transformation, my model highlights a different function that banks carry out: *risk transformation*. Moreover, the way the model is written is such that the risks are not separable. The bank has a single choice variable to manage two opposing risks. Therefore, it specifically highlights the potential for floating-rate loans to generate loan losses in response to contractionary monetary policy shocks.[30]

By differentiating equation (1.18) with respect to the monetary shock (ε), we can construct the model counterparts to the empirical impulse response functions:

$$\underbrace{\frac{\partial \Delta}{\partial \varepsilon}}_{\text{Profits IRF}} = \underbrace{f_L^* - \frac{D}{L}\beta^{dep}}_{\text{Net Interest Income IRF}} - \underbrace{f_L^* \theta'(\varepsilon)}_{\text{Provisions IRF}} \tag{1.19}$$

Equation (1.19) has a simple form. It states that impulse response function of profitability with respect to an interest rate shock is equal to the difference between the impulse response functions of net interest income and loan-loss provisions. Note that I have abstracted away from other components of bank income such as net noninterest income which include items such as fee income or salary expenses as these components are not core to understanding the loan-loss mechanism.

Importantly, equation (1.19) yields specific implications for the role of floating-rate loans.

[30] See Hellwig (1994) for a similar argument about the trade-off between interest rate risk and credit risk in relation to the Basel I regulatory framework.

First, let us look at the impulse response function of loan-loss provisions (the final term in (1.19)). We can see that the term is increasing in the share of floating-rate loans which tells us that loan losses will increase by more in response to a contractionary shock for banks with a higher floating share. Second, the impulse response function for net interest income also appears to be increasing in the floating share which tells us that net interest income will respond more positively in response to a contractionary shock for banks with a higher floating share.[31] This captures the trade-off between interest rate risk and credit risk described earlier. The model suggests that in response to a contractionary shock, banks with a higher floating share should experience a larger increase in net interest income but also a larger increase in loan-loss provisions.

The overall impact on bank profits will depend on the impact on net interest income relative to the impact on loan losses. However, we know from the aggregate data that profits fall, so one would expect that the impact of loan losses will dominate. In the next section, I test these implications using microdata.

1.6 Microdata Evidence

First, I aggregate the bank-level data to ensure it is reasonably close to the aggregate data series from the FDIC. The main variables that I am interested in exploring in this section are net interest income, provisions, and profits (all normalised by assets) as these are the core components of the model. In Figure 1.10 below, I show both the aggregate data from the FDIC and aggregated microdata from Drechsler et al. (2017) for each of these variables.

[31] Strictly speaking, it will also depend on the correlation between the share of floating-rate loans and the product of the deposit-loan ratio and deposit beta.

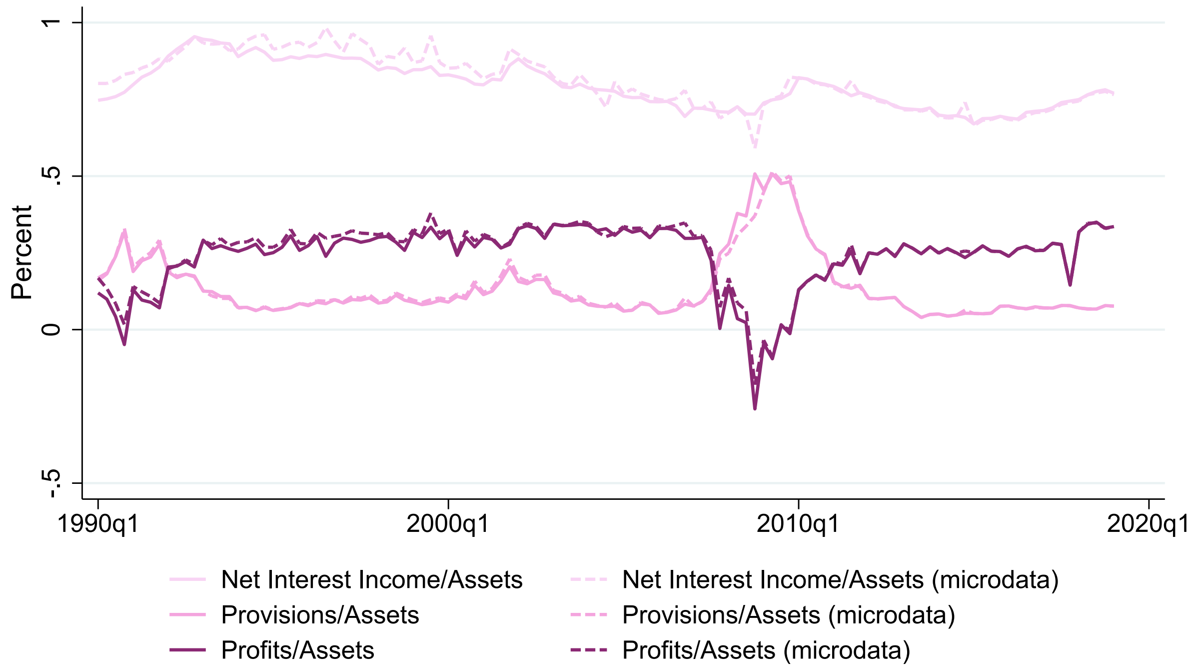

Figure 1.10. Comparing Aggregate Data to Aggregated Microdata

As can be seen from Figure 1.10, the microdata matches the macro data very well, albeit not perfectly. While there are some deviations in the mid-1990s and during the global financial crisis, both of these will not be in my estimation sample. The former will be excluded as data on the share of floating-rate loans begins in the late-1990s, while the latter is excluded, as in my earlier empirical analysis, due to the myriad changes to the regulatory architecture at the time. I define the floating share as follows:

$$f_L = \frac{\text{loans with repricing maturity of less than three months}}{\text{total loans}} \tag{1.20}$$

The numerator consists of two types of loans: floating-rate loans where the interest rate resets every three months (or more frequently) and fixed-rate loans with a remaining maturity of three months or less. While the latter group is not technically floating-rate loans, it will be considered as such for the purposes of my analysis. This is because loans that are fixed rate but with short maturity effectively act as floating-rate loans given they require frequent repricing.[32]

In Figure 1.11 below, I show the share of floating-rate loans in the time series and cross section. Figure 1.11a shows the time-series variation in the floating share for the aggregated banking sector. The aggregate floating share varies between 40% and 48%. Figure 1.11b takes the average floating share per bank over time and plots a histogram. As can be seen, there is considerable cross-sectional variation. While close to ten percent of banks have just under 20% of their loan portfolio composed of floating-rate loans, the distribution is clearly right-skewed.

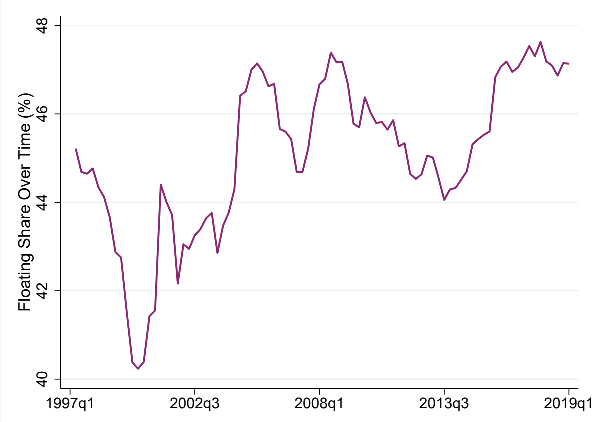

(a) Time-Series Variation in Floating Share

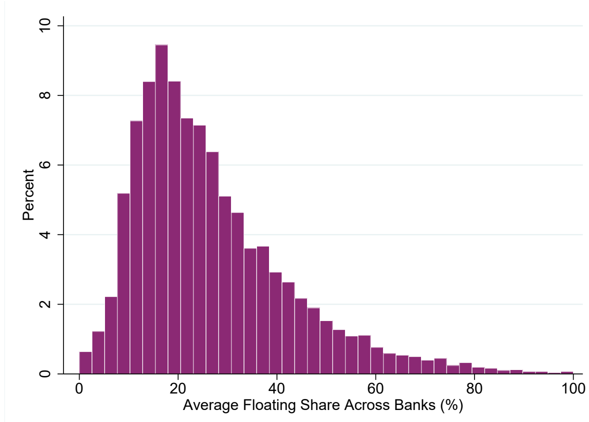

(b) Cross-Sectional Variation in Floating Share

Figure 1.11. The Share of Floating Rate Loans

While determining the specific causes of the floating share empirically is beyond the scope of this paper, it is worth documenting some of the average characteristics of banks with a lower average floating share relative to those with a higher average floating share within my estimation window (1999-2006). Specifically, I find that the average bank above the median floating share, relative to below the median, is substantially larger (over five times larger), has a higher share of commercial and industrial loans (20% versus 13%), has a slightly lower share of real estate-backed loans (62% versus 67%), has a lower share of personal loans (9% versus 14%), and is similarly profitable as measured by return assets (0.247% versus 0.250%).

[32] The main difference is in cases where a borrower is unable to refinance a fixed-rate loan of short maturity due to a high likelihood of default, but would have been forced to default on a floating-rate loan with a longer maturity. As such, my measure is likely to understate, rather than overstate, potential defaults.

Now that we have explored the floating-rate data, it is worth revisiting the model in the previous section. Recall that the model had a set of implications that depended on a bank's share of floating-rate loans. Specifically, the model suggests that in response to a contractionary monetary policy shock, banks with a higher floating share should experience higher net interest income but also higher loan-loss provisions, and that the impact on profits depends on the relative changes of the two components. I will test these implications using bank-level variation in the floating share. More precisely, I will estimate a panel local projection (a panel version of (1.1)) using data from 1999 to 2006 where the shock is interacted with the bank-specific floating share.[33] I also include horizon-specific bank fixed effects and consistent with (1.1), lags of the dependent variable and a quarter dummy. Therefore, I estimate the following specification for $h = 0 \dots 16$:

$$
\begin{aligned}
z_{i,t+h} = \alpha_{i,h} + \sum_{l=0}^{L} \beta_{h,l}^{(1)} Shock_{t-l} + \beta_{h}^{(2)} FloatShare_{i,t} + \sum_{l=0}^{L} \beta_{h,l}^{(3)} Shock_{t-l} \cdot FloatShare_{i,t} \\
+ \sum_{m=1}^{M} \gamma_{h,m} z_{t-m} + \sum_{q=2}^{4} \delta_q Quarter_{qt} + \varepsilon_{i,t+h} \qquad (1.21)
\end{aligned}
$$

Given the relatively short time series, I use four lags (i.e., $L = M = 4$). The main object of interest is the *interaction effect* $\{\beta_{h,0}^{(3)}\}_{h=0}^{H}$ for $h = 0 \dots 16$. A positive value of $\beta_{h,0}^{(3)}$ at horizon h implies that a higher floating share increases the response of $z_{i,t+h}$ to a monetary policy shock at time t. To ease interpretation and to document the magnitude, I will also show the *total effect* which is given by $\{\beta_{h,0}^{(1)} + \beta_{h,0}^{(3)} \cdot FloatShare_{i,t}\}_{h=0}^{H}$ for $h = 0 \dots 16$. The total effect measures the response of $z_{i,t+h}$ to a monetary policy shock at time t for specific values of the floating share. For illustrative purposes, I will show the 10th percentile and 90th percentile. However, these are only for illustrative purposes as the interaction effect directly captures the significance of the floating share for the responsive of z to the shock.

[33] I allow the floating share to be time-varying as from a macroeconomic perspective, it is useful to capture potential behavioural changes that result from the shock which might dampen its impact. However, I also estimate an alternative specification where I use the average floating share per bank which gives the same results.

Figure 1.12 below shows the response of net interest income to a contractionary monetary policy shock. Consistent with the model, the interaction effect in Figure 1.12a is mostly increasing in the floating share, albeit turning negative towards the end of the projection horizon. To better interpret the interaction effect, it is worth comparing 1.12b and 1.12c which capture banks with a low and high floating share, respectively. As expected, banks with a low floating share are more negatively impacted by a contractionary monetary policy shock.

Specifically, banks with a low floating share see a persistent fall in their net interest income. Intuitively, one can think of such banks as issuing largely fixed-rate loans and seeing the cost of their funding rise with interest rates. As such, their net interest income will fall. On the other hand, banks with a high floating share see their net interest income rise on impact and remain elevated for over two years as they generate more revenues on their loans, despite the cost of funding increasing. However, because these banks have passed on the interest rate risk to borrowers, borrowers eventually default which reduces loan repayments over time such that net interest income becomes negative, even for these high floating share banks. Overall, banks with a low floating share see a cumulative fall of around 0.4 percentage points in their net interest income while banks with a high floating share only experience a fall by about 0.1 percentage points. As the model suggests, banks with a high floating share are better hedged against interest rate risk.

(a) Interaction Effect

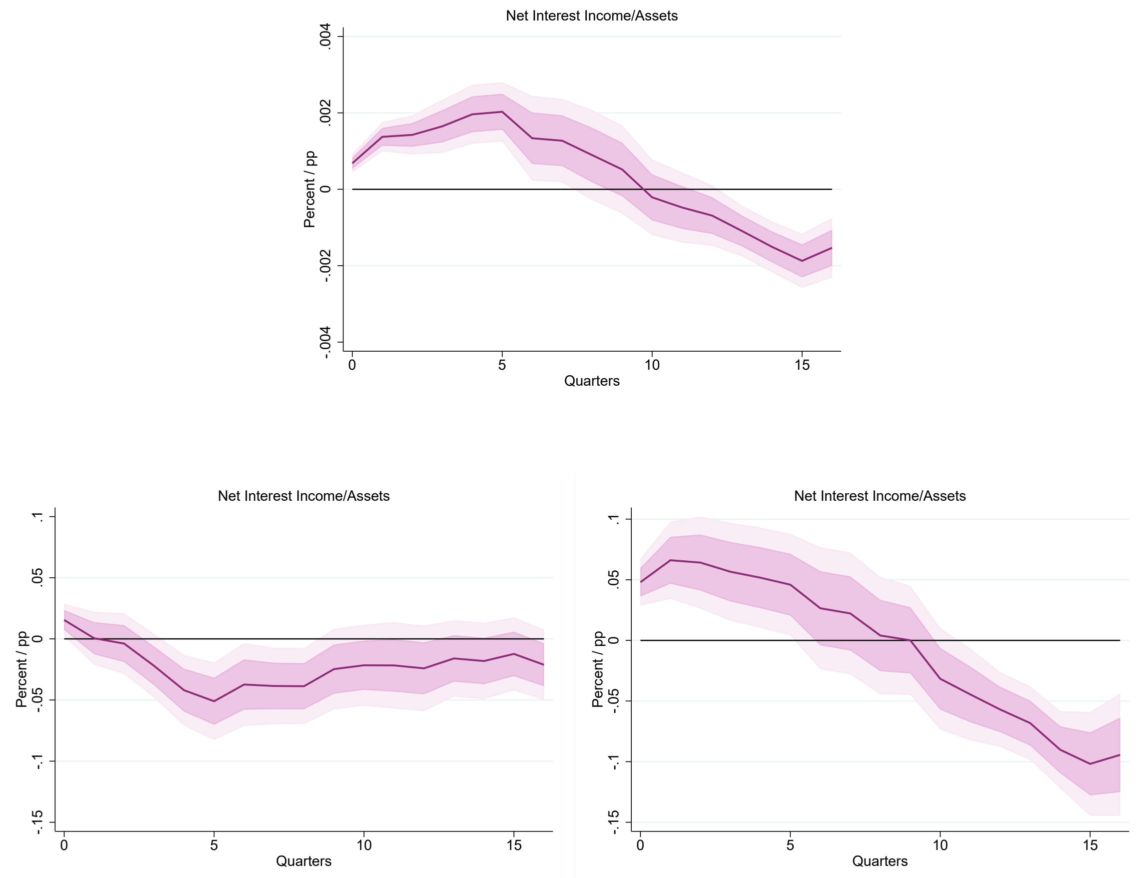

(b) Low Floating Share **(c)** High Floating Share

Figure 1.12. Net Interest Income Response To Contractionary MP Shock By Floating Share

In Figure 1.13, I repeat the analysis with loan-loss provisions instead of net interest income. Recall that the model predicts that loan losses should rise more for banks with a higher floating share than those with a lower floating share. Figure 1.13a plots the interaction effect which confirms this prediction of the model. Moreover, given the low floating share banks see a negative impact on their net interest income, one would expect minimal loan losses for this group as these banks do not appear to pass on their interest rate risk to their borrowers and so should not experience much loss from credit risk. Figure 1.13b shows precisely that low floating share banks see minimal loan losses. On the other hand, Figure 1.13c shows that loan losses rise significantly for banks with a high floating share, around three times as much as those with a low floating share. Indeed, this is specifically the trade-off emphasised by the model: banks are transforming (near-term) interest rate risk into (longer-term) credit risk.

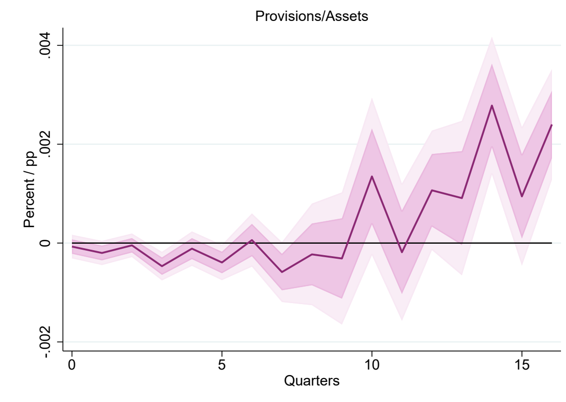

(a) Interaction Effect

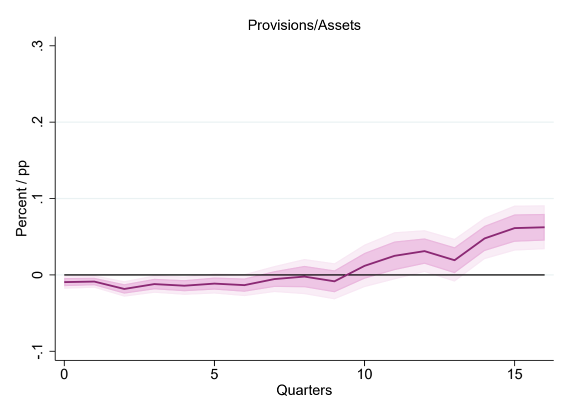

(b) Low Floating Share

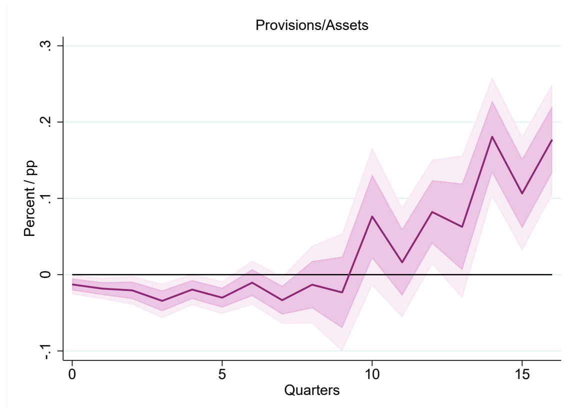

(c) High Floating Share

Figure 1.13. Loan-Loss Provision Response To Contractionary MP Shock By Floating Share

Finally, Figure 1.14 shows the same analysis but with overall bank profits. While the theoretical model in the previous section does not generate a directional prediction on overall profits, it does tell us that the impact on profits will be the difference between impulse response of net interest income and the impulse response of loan-loss provisions. One can see this immediately from Figures 1.14a, 1.14b, and 1.14c. For the low floating share banks (Figure 1.14b), profits are broadly flat, with a very small decline driven by the small increase in loan losses at the the end of the projection horizon. The high floating share banks (Figure 1.14c) initially see profits rise, driven by higher income on floating rate loans, but this hedged interest rate risk eventually becomes a crystallised credit risk. This results in a substantial rise in loan losses which leads to a significant overall decline in profits. Moreover, more of the variation in profits is due to loan losses than net interest income which is consistent with findings in the literature on the relative stability of bank net interest income (Drechsler et al. (2021)).

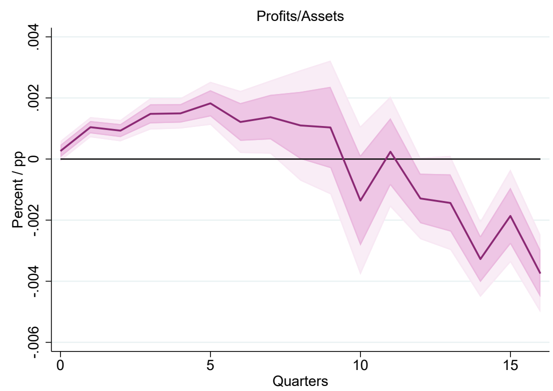

(a) Interaction Effect

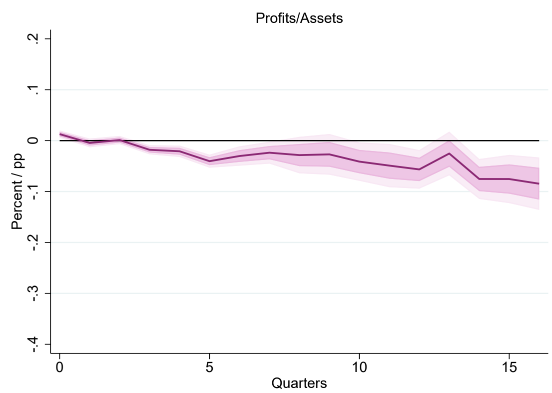

(b) Low Floating Share

(c) High Floating Share

Figure 1.14. Profit Response To Contractionary MP Shock By Floating Share

Taken together, Figures 1.12, 1.13, and 1.14 present evidence consistent with an important role for floating-rate loans in generating defaults, as in the theoretical model developed in Section 1.5. Put simply, the following story emerges. Banks are exposed to interest rate risk and as such issue floating-rate loans to hedge this risk. Given differential exposure to interest rate risk (e.g., through different deposit betas), banks issue different proportions of floating-rate loans. These floating-rate loans hedge the interest rate risk. However, because the interest rate risk is now with the borrower, it is now a credit risk for the bank. Ultimately, the credit risk component eventually becomes more important as loan losses offset the gain in net interest income, leading to a larger decline in profits for banks with a high floating share. The aggregate impact on profits is, as expected, in between the response of profits of low and high floating share banks.

One might ask why banks issue floating-rate loans if it leads to a decline in profits in response to a contractionary shock. First, it is worth noting that banks are not maximising value in order to just reduce the impact of interest rate shocks on profits. Second, and more intuitively, floating-rate loans are likely to see more benefit from expected interest rate changes rather than unexpected interest rate changes. While not explicitly modelled, as my focus is on the causal impact of monetary policy, a simple way to understand this point is the following. Higher interest rates lead to more income from floating-rate loans for banks. However, as I show, higher unexpected rates also result in less income for banks due to defaults. These defaults occur as borrowers are not well hedged against unexpected interest rate rises. The difference with higher expected interest rates is that expected interest rates are clearly not exogenous, they typically coincide with economic booms. As such, when interest rates rise, floating-rate borrowers experience higher income due to the economic boom, but also higher loan-servicing costs. This type of natural hedge is more pronounced with expected interest rate changes which are more procyclical than interest rate shocks. Therefore, while floating-rate loans may result in a decline in profits in response to contractionary monetary policy shocks, they are a far more effective hedge against typical interest changes.

Figure 1.15 below compares the response of loan-loss provisions to a contractionary monetary policy shock (1.15a) and to changes in the FFR (1.15b) where both result in a one percentage point rise in FFR.[34] I focus on comparing banks with high floating shares as they experience the largest increase in loan losses. The figures confirms that loan losses are substantially lower, and barely statistically significant, in response to changes in the FFR.

[34] Changes in the FFR are mostly expected changes in the interest rate but will include both expected and unexpected changes. As such, it should be considered an upper bound on the impact on loan losses.

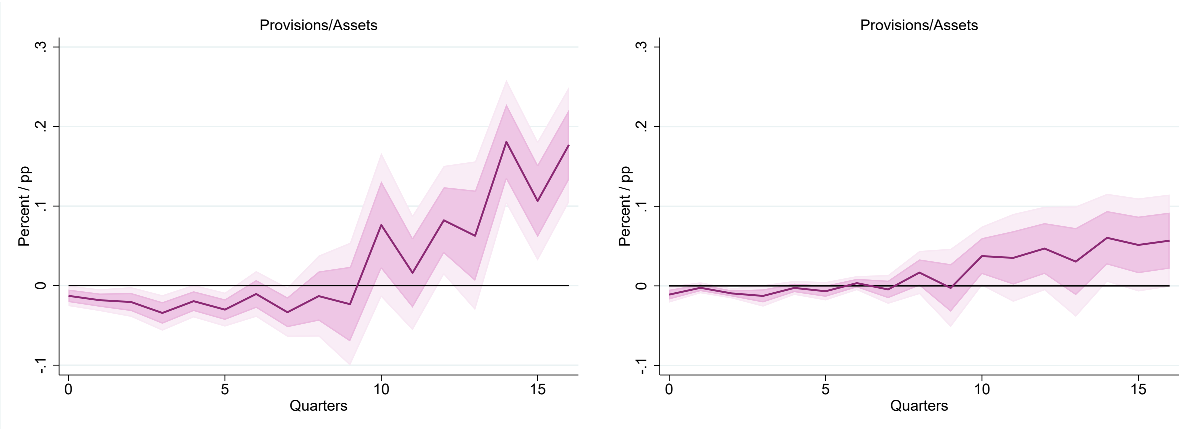

(a) Contractionary Monetary Policy Shock

(b) Higher Fed Funds Rate

Figure 1.15. Loan Loss Response to Higher Interest Rates for High Floating Share Banks

The fact that floating-rate loans, due to the pass-through of interest rates, generates an additional channel of defaults has been documented in different forms in the literature. For example, Campbell and Cocco (2015) find that when interest rates are high, defaults are higher for adjustable-rate mortgages relative to fixed-rate mortgages. In Europe, where floating-rate mortgages are more common, there is also evidence that floating-rate borrowers are more likely to default (Gaudêncio et al. (2019)).

In terms of external validity, my data stops before the GFC in 2007. Therefore, a natural question is whether the mechanism is likely to continue to apply. An important case study will be evaluating the consequences of the Fed's 2022 tightening cycle. While 2023Q3 banking data is yet to be released, there are early indications of rising defaults, in particular on floating-rate loans, which is precisely what my analysis would predict.[35] Moreover, one would potentially expect more significant loan losses in Europe where the share of floating-rate loans is higher. For example, in the UK, the Bank of England has projected that its rate hikes will lead to rising interest payments which will make it difficult for many companies to repay their debt.[36] More broadly, the implication of my analysis in this section is that the unintended consequences of

[35] https://www.ft.com/content/9a7e9746-516b-4d37-a966-97259ec8aca6

[36] https://www.bankofengland.co.uk/bank-overground/2023/how-vulnerable-are-uk-companies-to-higher-interest-rates

higher interest rates on the stability of the banking sector are potentially more severe where the share of floating-rate loans is higher.

1.7 Conclusion

In this paper, I explore the following question: do contractionary monetary policy shocks make banks safer through reducing their leverage? While a vast theoretical literature claims the answer is yes, I show empirically that the answer is actually no. Not only is raising interest rates ineffective in reducing bank leverage, it is actively counterproductive as it increases leverage instead. I show this result is robust to varying specifications and using different measures of monetary policy shocks.

Next, I show empirically why leverage rises in response to contractionary monetary policy shocks. Higher interest rates increase loan losses for banks. This reduces bank profits overall which subsequently reduces bank equity. The fall in equity drives an increase in bank leverage. I term this mechanism the loan-loss mechanism. Moreover, I show empirically that the loan-loss mechanism can explain nearly all the variation in bank leverage in response to monetary policy shocks. Finally, I show that while loan losses and leverage increase in response to monetary policy shocks (where the FFR rises), loan losses do not rise and leverage falls in response to contractionary oil shocks (where the FFR does not rise). This analysis provides suggestive evidence, at the aggregate level, of the importance of the rise in the FFR specifically and hence floating-rate loans. This highlights the importance of understanding bank balance sheets, and in particular the structure of the loan portfolio, in order to understand the transmission of monetary policy.

Moving on to the theoretical literature, I show that the divergence between the theoretical claims and empirical evidence is largely a result of three broad modelling choices and that there is one important factor that can help rectify this. The first modelling choice relates to models that rely on profitability rising in response to a contractionary monetary policy shock which is inconsistent with the empirical evidence. The second relates to models that incorrectly rely on the procyclicality of bank leverage and so erroneously conclude that leverage declines in response to rising interest rates. The third relates to models that rely on the substitution effect through which higher rates reduce bank leverage which is inconsistent with the observed evidence. The crucial missing factor in this eclectic mix of models is a loan-loss mechanism that is connected to the share of floating-rate loans issued by a bank.

I develop a banking model that emphasises the role of floating-rate loans and credit risk. Banks optimise by choosing the floating share of their loan portfolio which acts as a hedge against interest rate risk, but generates credit risk for the bank as the interest rate risk is now held by borrowers. A key insight of the model is that banks are doing risk transformation, and that this implies a trade-off between managing interest rate risk and credit risk. The model predicts that banks with a higher share of floating-rate loans will see greater loan losses in response to a contractionary monetary policy shock. I confirm this prediction using microdata, specifically, bank-level variation in the floating share.

My results have important implications for using monetary policy for financial stability purposes. First, an important reason to support a monetary policy strategy that targets financial stability (i.e., 'leaning against the wind') is the claim that higher rates reduce bank leverage. In this paper, I have shown this claim to be empirically false. Therefore, this paper lends support to the conclusions of Bernanke (2015) and Svensson (2017) that monetary policy should focus on its mandate of price stability, leaving issues of financial stability to macroprudential policy. However, my results also suggest that floating-rate loans are one specific way through which monetary policy creates unintended vulnerabilities in the banking sector. This is particularly pertinent in economies with a greater share of floating-rate loans (e.g., Europe). Indeed, future research could consider this novel trade-off for monetary policy: a higher share of floating-rate loans can increase the potency of monetary policy (e.g., Calza et al. (2013) and Auclert (2019)) but comes at the cost of a more vulnerable financial sector, as documented in this paper.

1.8 Chapter Acknowledgements

Chapter 1 is currently being prepared for submission for publication of the material. The dissertation author was the sole investigator and author of this paper.

Chapter 2

Do Central Bank Cycles Drive Stock Returns? New Evidence From The US, UK, And Japan

2.1 Introduction

If you invested \$100 in the stock market at the start of 1994, you would have \$768 in 2016. However, if you only held stocks in even weeks of the Federal Open Markets Committee (FOMC) cycle[1] over the same time period, your \$100 would become a striking \$1,522. This ‘even-week’ result was meticulously derived by Cieslak et al. (2019) (hereafter CMVJ), inspired by the earlier work of Lucca and Moench (2015). In this paper, I evaluate the robustness of the even-week result as well as the underlying mechanism proposed by CMVJ.

The question of whether the Federal Reserve (the Fed) has a substantial impact on the stock market, as well as how it does, are important questions and have meaningful implications for our understanding of monetary policy and the functioning of central banks. However, the question is fraught with challenges given the difficulty in identifying monetary policy shocks (Ramey (2016)). Many papers seek identification by exploring the impact of monetary policy around the time of central bank announcements (see for example Lucca and Moench (2015) and Brusa et al. (2020)).

[1] Even weeks refer to weeks 0, 2, 4, and 6 of the FOMC cycle. The FOMC cycle captures all the days from one FOMC announcement to the next. See Table 2.1 for a full mapping of cycle days to cycle weeks.

CMVJ use neither esoteric measures of monetary policy shocks nor the high-frequency approach around announcements. Instead, they focus on the evolution of stock returns over the full cycle of days between scheduled FOMC meetings and find that the equity premium is earned entirely in even weeks of the FOMC cycle, claiming that the FOMC cycle in stock returns appears to be a general phenomenon that is strengthening over time. By ruling out alternative explanations methodically, they claim that the even-week result is *causally* driven by systematic informal communication around biweekly board meetings at the Fed. CMVJ's findings are significant and novel. Indeed, not only have they sparked an important debate on the costs and benefits of informal communication (leaks), they have cautioned the Fed, arguing that "any such benefits must be balanced against the risk of insider trading and informal communication undermining the public's trust in financial markets and the Fed."

Moreover, the novelty and significance of CMVJ's result has both driven momentum into this area of the literature (see, for example, Finer (2018), Laarits (2020), Hu et al. (2022), Bradley et al. (2023)) and their findings have been widely cited in the broader central bank communication literature (Ai and Bansal (2018), Bianchi et al. (2023), Gorodnichenko et al. (2023), Masciandaro et al. (2024)). However, CMVJ's findings are not only influential in central banks and academia, they have also been widely reported in the media, with both The Economist[2] and The Wall Street Journal[3] featuring articles discussing the even-week result. Lastly, CMVJ show that their findings are also relevant globally. They find that the FOMC cycle, in addition to driving US stock returns, also drives international stock returns.[4]

In the context of the findings above, this paper has three key contributions. First, I find that the even-week result does not hold out-of-sample, casting doubt on the claim that it is a general phenomenon. Furthermore, while CMVJ show their result for the period 1994-2016, I show that when using data up to the end of 2023, the even-week result loses statistical significance

[2] https://www.economist.com/finance-and-economics/2016/09/03/the-long-arm-of-the-fed

[3] https://on.wsj.com/3li9ypk

[4] Relatedly, Brusa et al. (2020) show that high equity returns around monetary policy announcements are unique to the Fed.

as early as 2004, with the even-week coefficient trending down following the financial crisis. I explain this by documenting that biweekly board meetings, which are key to CMVJ's proposed mechanism, stopped being biweekly from 2004 onwards. Second, I show that in the much narrower sample of 1994-2003, where the even-week result holds and board meetings are at least biweekly, there appear to be a few important days (rather than weeks) of the FOMC cycle. These days have excess stock returns that are outliers and driven by factors unrelated to the FOMC cycle. Once removing these outliers, I show that even in the 1994-2003 period, the even-week result no longer holds. Finally, in the context of the FOMC cycle driving international stock returns, I show that neither UK nor Japanese stock returns are driven by FOMC cycle when one controls for a pre-announcement effect. Taken together, my results suggest that the even-week result may have simply been a statistical artifact, rather than driven by systematic leaks every two weeks by the Fed for over two decades.

The remainder of this paper is as follows. Section 2.2 describes the data. Section 2.3 evaluates the robustness of the even-week result. Section 2.4 inspects the proposed mechanism. Section 2.5 considers the even-week result in the UK and Japan. Section 2.6 provides a conclusion.

2.2 Data

A key contribution of this paper is the collection of a variety of datasets, transformed in a way to ensure consistency and comparability across central bank cycles of varying lengths. This section is split into data relating to central bank meetings, financial data, and Fed Chair calendar data. While much of my data goes up to 2016, in line with that of CMVJ, the extended data for the US goes up to the end of 2023 which expands CMVJ's original sample by seven years.

2.2.1 Central Bank Meeting Data

I collect data on central bank meetings for the US, the UK, and Japan. I then compute days in terms of the central bank's cycle time. I use the definition of central bank cycle time as in CMVJ.

The cycle time captures the rate-setting process of the central bank. CMVJ define the FOMC cycle by having week 0 of the FOMC cycle start the day before a scheduled FOMC announcement. The announcement day is considered day 0, therefore week 0 starts on day -1. Their rationale for having week 0 start the day before the announcement rather than the day of the announcement is to capture the pre-announcement effect documented by Lucca and Moench (2015) (i.e., large excess stock returns in advance of FOMC meetings), though one could argue this is distinct from the proposed mechanism in CMVJ (see Section 2.4). The definition of cycle time is documented in the table below.

Table 2.1. Defining weeks of central bank cycles

Week of the cycle	Days counted under Cieslak et al. (2019)
-1	$-6,\ldots,-2$
0	$-1,\ldots,3$
1	$4,\ldots,8$
2	$9,\ldots,13$
3	$14,\ldots,18$
4	$19,\ldots,23$
5	$24,\ldots,28$
6	$29,\ldots,33$

Fed Meetings

I start with the data compiled by CMVJ who collects FOMC meeting dates from 1982 to 2016. I then update their data to include FOMC meetings up to the end of 2023, using the Federal Reserve website. Figure 2.1 below shows the days on which the meetings took place for the post-1994 period as this will be the period relevant for analysis.[5] There are 216 FOMC meetings in total and as can be seen by the peaks in Figure 2.1 below, the FOMC meets eight times per year.

[5] The Fed only started publicly announcing its decision following a scheduled FOMC meeting in 1994, so interpretations before this date are less meaningful.

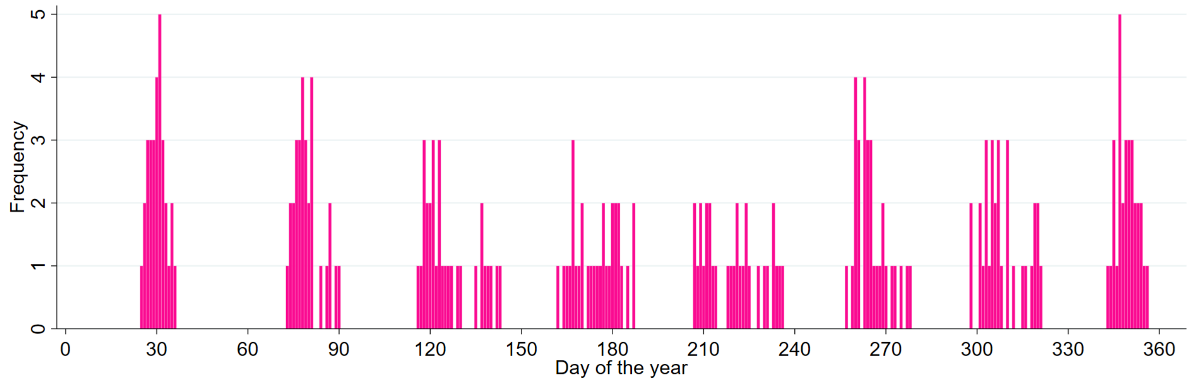

Figure 2.1. FOMC Meeting Frequency and Timing

Given the importance of the Fed's board meetings for CMVJ's argument, I also update CMVJ's data to include board meetings up to the end of 2023. These are discussed in greater detail in Section 2.4.1.

Bank of England (BoE) Meetings

Using the BoE website, I collect information on meetings of the Monetary Policy Committee (I will refer to these as BoE meetings). My sample starts in July 1997 (the BoE gained operational independence in June 1997). Unlike the FOMC, the BoE meets monthly. However, in late 2016, it changed its meeting schedule to an eight-meetings-a-year schedule. Therefore, for consistency, I focus on the period from 1997 to 2016. There are 232 BoE meetings in total and we can see the 12 monthly meetings in Figure 2.2 below.

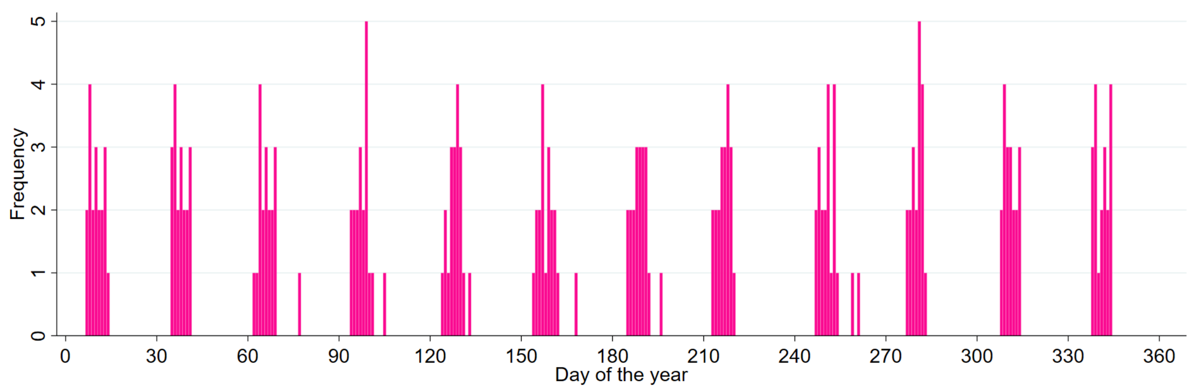

Figure 2.2. BoE Meeting Frequency and Timing

Bank of Japan (BoJ) Meetings

I collect information on BoJ monetary policy meetings from the BoJ website. Given the BoJ gained independence in April 1998, I start my sample in May 1998. The BoJ's meeting schedule was much less regular than that of either the FOMC or BoE (see Figure 2.3). In line with CMVJ, I exclude unscheduled meetings. As highlighted by Brusa et al. (2020), up until 2005, the BoJ progressively decreased the number of meetings from 20 per year, before settling on 14 in 2006. Like the BoE, the BoJ decided to shift to an eight-meeting schedule in 2016. I focus on the period from 1998 to 2016 which most closely matches the BoE sample and the FOMC sample of CMVJ. This sample consists of 286 BoJ meetings.

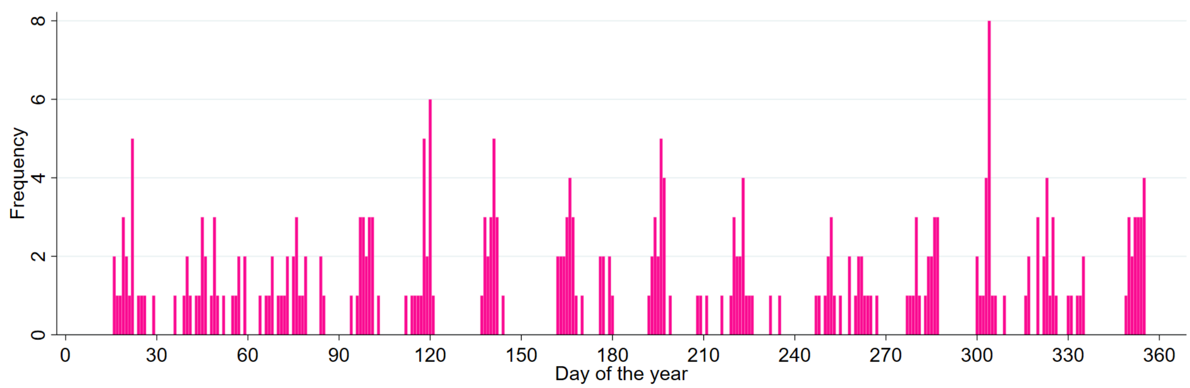

Figure 2.3. BoJ Meeting Frequency and Timing

2.2.2 Financial Data

Data on the Fed funds rate is from FRED while data on Fed funds futures contracts and Eurodollar futures contracts is from Bloomberg.

US and Japanese daily excess return data comes from Kenneth French's website while the equivalent for the UK comes from Gregory et al. (2013) with updates from the Xfi Centre for Finance and Investment at the University of Exeter Business School. Table 2.2 below shows the key summary statistics for these variables.

Table 2.2. Excess Return Summary Statistics

Variable	Obs	Mean	Std. Dev.	First-Order AC
US Daily Excess Return	7,821	0.04	1.17	-0.07
UK Daily Excess Return	4,864	0.02	1.15	-0.01
Japan Daily Excess Return	4,957	0.02	1.42	-0.06

2.3 Robustness of the Even-Week Result in the US

This section explores the robustness of the even-week result documented by CMVJ. First, I replicate CMVJ's main results and test their robustness out-of-sample. Then, I test subsample stability by doing an expanding window regression and splitting the thirty years of data into three ten-year samples.

2.3.1 Out-of-Sample Robustness

CMVJ show that their results are robust out-of-sample. Specifically, they split their sample into two periods: 1994-2013 and 2014-2016. The latter three-year period constitutes an out-of-sample test for CMVJ as the initial draft of their paper only employed data from 1994 to 2013. Not only do they show that their results hold in their out-of-sample period, leading them to conclude that the FOMC cycle in stock returns is a general phenomenon, they further claim that the even-week result is strengthening over time. I replicate CMVJ's results and test whether the even-week result remains robust to a more recent out-of-sample period. My out-of-sample period is the seven-year period from the start of 2017 to the end of 2023. Further, I combine this with CMVJ's out-of-sample period so that I can also test whether the result holds for the ten-year period from the start of 2014 to the end of 2023. Here I obtain my first contribution in this paper: the FOMC cycle in stock returns no longer holds out-of-sample in either my out-of-sample or the combined out-of-sample.

Table 2.3 reports regressions of daily excess US stock returns on FOMC cycle dummies. t- statistics are calculated on the basis of robust standard errors as is the case in the rest of this paper. Columns 1 and 2 of Table 2.3 replicate the results reported by CMVJ in Table I Panel B of their paper. The interpretation of the results is that, between 1994 and 2013, the average excess return per day is 13.6 bps higher on days that fall in week 0 in FOMC cycle time and 9.9 bps higher on days that fall in week 2, 4, or 6 compared to days that fall in odd weeks in FOMC cycle time. The coefficient on the week 2, 4, or 6 dummy strengthens substantially in the 2014-2016 period. However, column 3 shows that the result does not hold in the 2017-2021 period.[6] The week 0 dummy, although insignificant, is now negative rather than positive while the dummy for even weeks 2, 4, or 6 is indistinguishable from zero. Note that the number of observations in CMVJ's out-of-sample period (783) is much smaller than in mine (1824). Finally, in column 4, I combine the two out-of-sample periods so that there is one larger out-of-sample test with ten years of data from the start of 2014 to the end of 2023. I find that none of the coefficients are statistically significant.

Table 2.3. Regressions of Daily Excess U.S. Stock Returns on FOMC Cycle Dummies

	CMV Main 1994 to 2013	CMV OoS 2014 to 2016	New OoS 2017 to 2023	Combined OoS 2014 to 2023
Dummy=1 in Week 0	0.136***	0.174*	-0.139	-0.0450
	(2.76)	(1.92)	(-1.58)	(-0.67)
Dummy=1 in Week 2,4,6	0.0993***	0.176***	-0.00773	0.0475
	(2.65)	(2.67)	(-0.13)	(1.02)
N	5214	783	1824	2607

t statistics in parentheses
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

While Table 2.3 replicates and extends CMVJ's out-of-sample test, it does not test the significance of each individual even-week dummy (i.e., dummy=1 in week 0, 2, 4, 6). Table 2.4 therefore repeats the exercise in Table 2.3 but with individual even-week dummies. A few points are worth noting. First, the even-week result appears less significant in CMVJ's out-of-sample period (2014-2016) when using individual even-week dummies. In my out-of-sample period

[6] My results remain consistent when using the out-of-sample period 2017-2019, i.e., excluding 2020 onwards given potential concerns about the global pandemic.

(2017-2023) and the combined out-of-sample period (2014-2023), the results are consistent with those in Table 2.3: the FOMC cycle in stock returns does not hold out-of-sample.

Table 2.4. Regressions of Daily Excess U.S. Stock Returns on FOMC Cycle Dummies

	CMV Main 1994 to 2013	CMV OoS 2014 to 2016	New OoS 2017 to 2023	Combined OoS 2014 to 2023
Dummy=1 in Week 0	0.136*** (2.76)	0.174* (1.92)	-0.139 (-1.58)	-0.0450 (-0.67)
Dummy=1 in Week 2	0.0811* (1.70)	0.146* (1.82)	0.0243 (0.31)	0.0610 (1.01)
Dummy=1 in Week 4	0.107** (1.99)	0.200** (2.41)	-0.0444 (-0.60)	0.0291 (0.51)
Dummy=1 in Week 6	0.177** (1.98)	0.325 (0.52)	0.0365 (0.25)	0.104 (0.62)
N	5214	783	1824	2607

t statistics in parentheses
* p<0.10, ** p<0.05, *** p<0.01

The lack of robustness of the even-week result naturally leads to the question of when it stopped being robust. One puzzle when reading the final version of CMVJ's paper (published in 2019) is that given the initial paper was available in 2014, why had the even-week result not been arbitraged away? Indeed, looking at column 4 of Table 2.3 and Table 2.4, it would appear the result is no longer significant after the initial draft was made available.[7] This would be consistent with the findings of McLean and Pontiff (2016) who find that investors learn about mispricing from academic publications.

2.3.2 Subsample Stability

In order to understand when the even-week result began to lose significance, I estimate a regression of the daily excess stock return on the even-week dummy (i.e., dummy=1 in week

[7] The first draft of the CMVJ paper was available on 23 April 2014. See https://faculty.haas.berkeley.edu/morse/research/papers/cycle_paper_cieslak_morse_vissingjorgensen.pdf

0, 2, 4, or 6) using an expanding window going backwards from 2023. First, I do the analysis using CMVJ's time horizon (i.e., 1994-2016) as shown in Figure 2.4. Then, I use my updated time horizon with data going up until the end of 2023 as shown in Figure 2.5.[8] Each point in the figure represents the even-week regression coefficient estimated for the time horizon along the horizontal axis. For example, the last marker in Figure 2.4 represents the excess return when estimating the even-week coefficient using data from the start of 2015 to the end of 2016.

Figure 2.4 clearly supports the CMVJ interpretation. Specifically, that the even-week effect is robust in their sample period (1994-2016) and it appears to be strengthening over time. CMVJ argue that "the FOMC cycle in stock returns appears to be a general phenomenon, present since 1982, but strengthening over time in economic magnitude" (p.2208). However, Figure 2.5 tells a surprisingly different story. First, we can see that the result loses most of its statistical significance consistently from as early as 2004 when it is only statistically significant at the ten percent level. Moreover, for any of the samples from 2009 onwards, the result is not statistically significant at all (See Table B.1 in Appendix B.1 for a table with the specific coefficients and *t*-statistics used in Figure 2.5). Second, the even-week effect goes from being positive (as in CMVJ) to negative, with the decline starting after the global financial crisis.[9] Together, these two points suggest that CMVJ's result no longer holds and is not as robust when using a longer time horizon.

[8] The results are similar when excluding the global pandemic, i.e., ending the sample at end-2019.

[9] Note that in the last two years, the coefficient is again positive but extremely noisy.

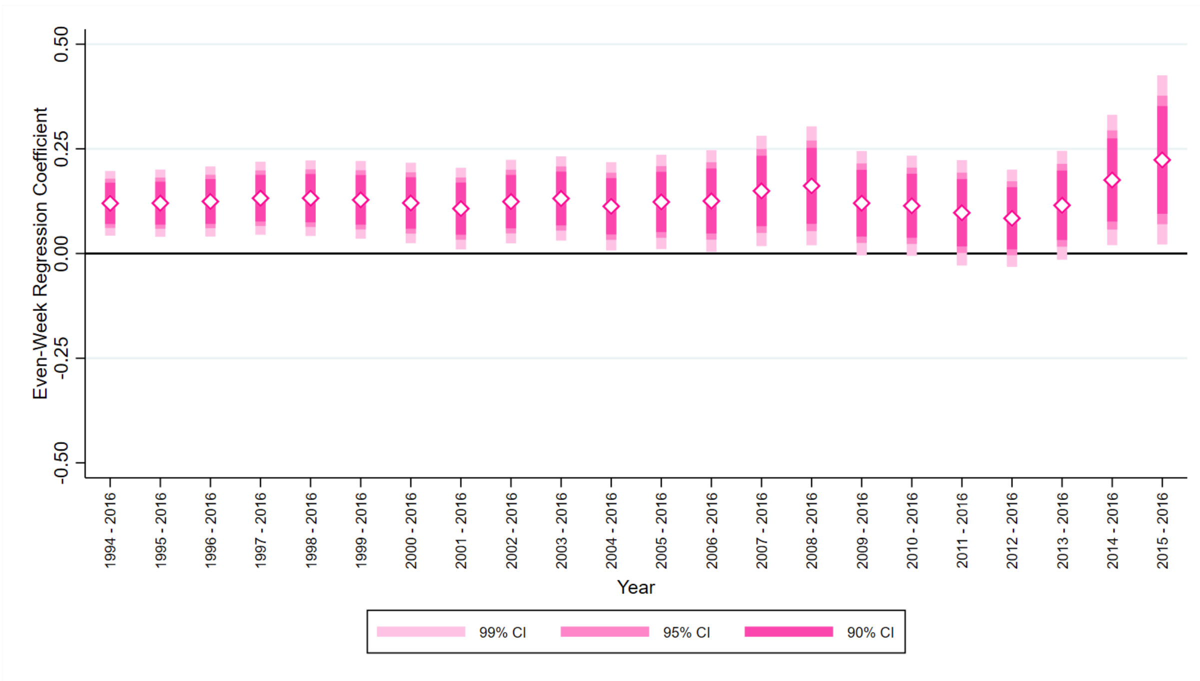

Figure 2.4. Expanding Window Regression, 1994-2016

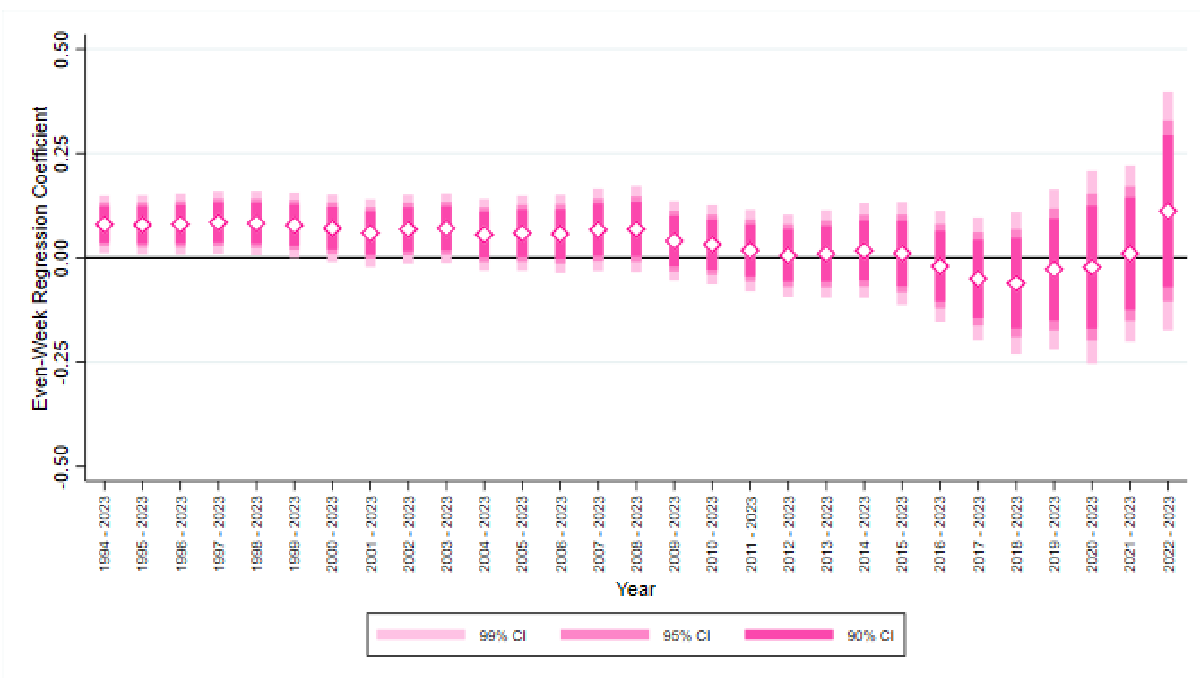

Figure 2.5. Expanding Window Regression, 1994-2023

While the expanding window regressions suggest a lack of robustness when varying the time horizon, the number of observations is strictly decreasing over time which could lead to larger standard errors. Therefore, I will now examine subsample stability by splitting the thirty years of data into three ten-year periods: 1984-2003, 2004-2013, and 2014-2023. As in the previous section, I show the results when separating out the week 0 dummy from the week 2,4,6 dummy given the former could be more related to the already-documented FOMC drift (Lucca and Moench (2015)) as well as only individual even-week dummies.

Tables 2.5 and 2.6 present these results. As can be seen, the CMVJ result is only statistically significant for ten of the thirty years. Specifically, it only appears to hold consistently for the 1994-2003 period. Given the sample is smaller, one would expect potentially larger standard errors, but all the statistically significant coefficients also shrink substantially relative to the 1994-2003 period. Interestingly, the 2004-2013 period which is entirely contained within CMVJ's original sample is not statistically significant at even the ten percent level. Finally, in the 2014-2023 sample, the results remain statistically insignificant and the coefficients are again mostly smaller, and in some cases even negative.

Table 2.5. Subsample Stability

	1994 to 2003	2004 to 2013	2014 to 2023
Dummy=1 in Week 0	0.149**	0.123	-0.0450
	(2.47)	(1.57)	(-0.67)
Dummy=1 in Week 2,4,6	0.119**	0.0795	0.0475
	(2.35)	(1.44)	(1.02)
N	2608	2606	2607

t statistics in parentheses

* p<0.10, ** p<0.05, *** p<0.01

Table 2.6. Subsample Stability

	1994 to 2003	2004 to 2013	2014 to 2023
Dummy=1 in Week 0	0.149**	0.123	-0.0450
	(2.47)	(1.57)	(-0.67)
Dummy=1 in Week 2	0.0788	0.0835	0.0610
	(1.18)	(1.23)	(1.01)
Dummy=1 in Week 4	0.151**	0.0674	0.0291
	(2.10)	(0.85)	(0.51)
Dummy=1 in Week 6	0.205*	0.143	0.104
	(1.70)	(1.07)	(0.62)
N	2608	2606	2607

t statistics in parentheses
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

The lack of robustness of the week 0 dummy is likely to reflect an already documented finding: the FOMC drift has been disappearing (Kurov et al. (2021)). However, the lack of robustness of the other even-week dummies challenges the robustness of the CMVJ result not only out-of-sample, but also potentially in-sample. Furthermore, it is unlikely that investors learned about the mispricing and subsequently arbitraged away the even-week excess returns given that the result stops holding as early as 2004 while the initial draft of the CMVJ paper was released a decade later, in April 2014. Therefore, to understand why the CMVJ result appears to stop holding as early as 2004, one must inspect the proposed mechanisms.

2.4 Inspecting the Mechanism

The key mechanism underpinning the CMVJ result is the following. Board members meet during even weeks of the FOMC cycle. There is an information exchange and update during these meetings, and according to CMVJ, such meetings result in information being (informally) leaked. Therefore, they argue, excess stock returns in even weeks are caused by leaks following biweekly board member meetings.[10] In this section, I examine the role of board meetings as well

[10] CMVJ also note that there was unexpectedly accomodative monetary policy during their sample which may have contributed to the excess returns. However, I show that monetary policy is similarly accomodative over the

as whether other evidence is consistent with CMVJ's mechanism.

2.4.1 Role of Board meetings

If CMVJ's mechanism holds, one would expect that board meetings are held at least biweekly throughout the year, which would entail at least 26 annual meetings.[11] If there are significantly fewer than 26 annual board meetings, it would be difficult for CMVJ's mechanism to be true as the board members could not be meeting every two weeks, which is needed for the even-week result. To examine the role of board meetings, I first document the number of annual board meetings in Figure 2.6 below.

three ten-year periods between 1994 and 2023. Specifically, in Appendix B.2, I look at the monetary policy shock series of Bu et al. (2021) as it covers both the CMVJ sample as well as most of my extended sample (it stops in September 2023). I show in Figure B.1 that there is little difference between the average shock over the three ten-year periods between 1994 and 2023. In fact, the average shock over the 1994-2003 period is less expansionary than over the 2004-2013 period, when the even-week result no longer holds.

[11] Note that there could be more than 26 meetings. This is not inconsistent with CMVJ's claim as their claim rests on biweekly meetings having a special role. They argue that those specific meetings are when information is updated and exchanged.

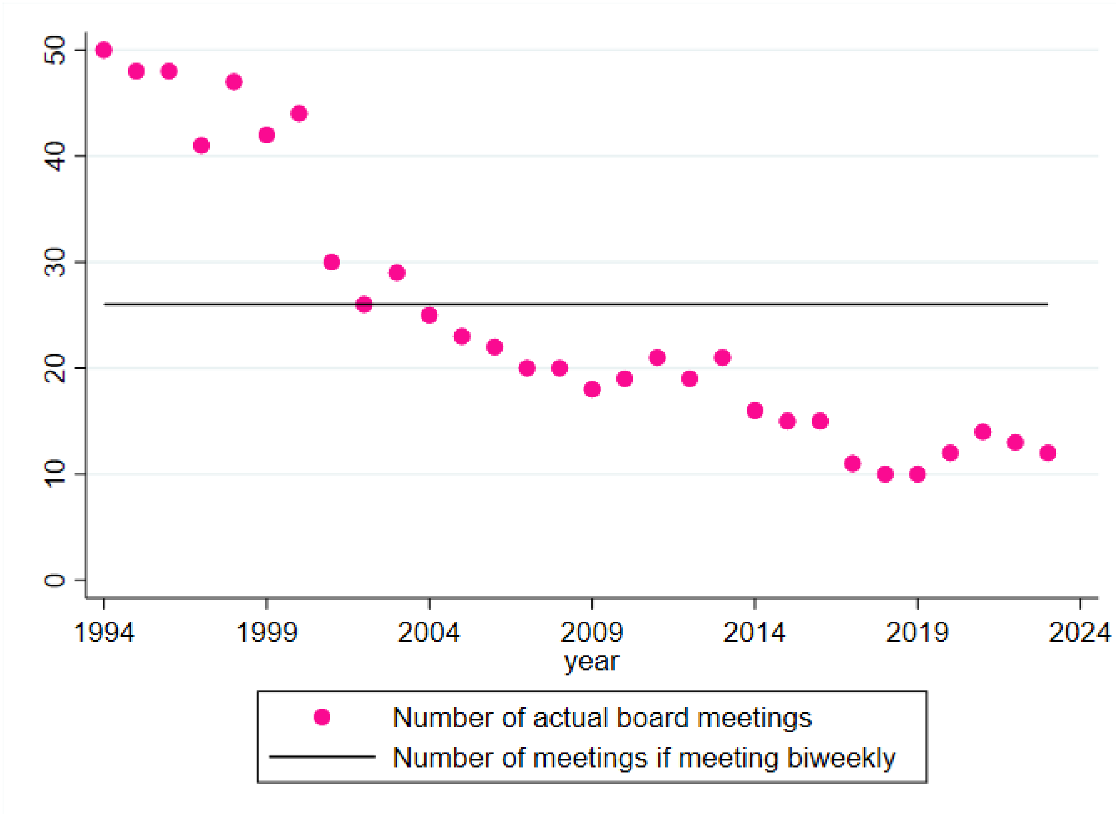

Figure 2.6. Number of Board Meetings per year

As can be seen in Figure 2.6, the annual number of board meetings has been steadily declining. The figure shows a horizontal line at 26 annual meetings which represents biweekly meetings. Interestingly, the number of annual meetings fall below 26 from 2004 onwards. This is precisely when the CMVJ results appears to no longer hold as shown earlier in Tables 2.5 and 2.6. For example, years 2014, 2015, 2016 are all within CMVJ's main sample. However, each of these years has no more than 16 annual meetings which implies a meeting frequency of less than one every three weeks. This appears to be inconsistent with the mechanism that even-week excess returns are causally driven by information leaks following biweekly board meetings.

While Figure 2.6 casts doubt on CMVJ's proposed mechanism, it does not provide conclusive evidence against the mechanism. For example, it could be the case that some meetings were biweekly which lead to excess returns and others were at a much lower frequency, such that the average meeting frequency is not biweekly. In order to test the role of board meetings more precisely, I replicate the regression of CMVJ by regressing excess returns on interaction terms between each of the even-week dummies and a dummy for whether any of days $t-5$ to $t-1$ had a board meeting. A positive value would indicate that the excess returns specifically followed board meetings in even weeks and provide the clearest support for CMVJ's mechanism.[12] As in my prior regressions, I examine these interactions for three periods: 1994-2003, 2004-2013, and 2014-2023. These periods also coincide with changes in the frequency of board meetings in Figure 2.6.

Table 2.7 below documents my results. Consistent with the results showing that the even-week effect does not hold from 2004 onwards, columns 2 (2004-2013) and 3 (2014-2023) show the interaction between board meeting dummy and even-week dummy is statically insignificant. Moreover, the coefficients are substantially smaller than those in column 1 (1994-2003).[13] Therefore, columns 2 and 3, when combined with the results presented earlier, provide evidence that the even-week result and the supporting mechanism did not hold, at least since 2004. While column 1 provides evidence consistent with CMVJ's mechanism, if one considers that CMVJ's entire sample is 1994-2016, then this result holds in less than half the sample.

[12] Note, in line with CMVJ, I also include two further interaction terms: the interaction between the board meeting dummy and a dummy for even weeks with no board meeting as well as the interaction between the board meeting dummy and an odd-week dummy.

[13] As mentioned before, the week 0 dummy, while barely significant, is likely to capture the announcement effect rather than represent the biweekly board meetings.

Table 2.7. Even-Week Effect and Board Meetings

	1994-2003	2004-2013	2014-2023
[Week 0]*[Board meeting in preceding 5 days]	0.274***	0.152*	0.0804
	(3.47)	(1.84)	(0.93)
[Week 2]*[Board meeting in preceding 5 days]	0.273***	0.152	0.0747
	(3.08)	(1.64)	(0.67)
[Week 4]*[Board meeting in preceding 5 days]	0.294***	0.0749	0.0654
	(3.16)	(0.60)	(0.75)
[Week 6]*[Board meeting in preceding 5 days]	0.266*	0.135	-0.267***
	(1.83)	(0.70)	(-7.09)
[Even-week]*[No board meeting in preceding 5 days]	0.0918	-0.00861	-0.00438
	(0.89)	(-0.12)	(-0.08)
[Odd-week]*[Board meeting in preceding 5 days]	0.165**	-0.0790	-0.00767
	(2.48)	(-1.09)	(-0.11)
N	2608	2606	2607

Dummy variables in square brackets, *t* statistics in parentheses
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

While column 1 of Table 2.7 provides evidence in support of the role of board meetings from 1994-2003, it is not conclusive. For example, it is possible that board meetings coincided with excess returns in even weeks for reasons unrelated to information leaks by the Fed. As such it is worth exploring whether other evidence is consistent with CMVJ's proposed mechanism that leaks about monetary policy following board meetings caused excess stock returns in even weeks.

2.4.2 Impact on Fixed Income

Given CMVJ's argument is that leaks following biweekly board meetings systematically push up stock prices in even weeks, one would expect fixed income assets to also display an even-week pattern. Indeed, one might expect the response of fixed income to be even stronger than that of stocks as they are arguably much more sensitive to news about monetary policy. While CMVJ explore the impact on some fixed income assets, in this section, I will document the robustness of their result to a wider array of fixed income assets. Specifically, I examine the responses of 1-year, 2-year, and 10-year US Treasuries, the second and fourth Fed Funds futures contracts, and the fourth Eurodollar futures contract. I will focus specifically on the period 1994-2003 as I have already shown their even-week result and underlying mechanism does not appear to hold from 2004 onwards.

In Table 2.8 below, I document the daily change in US Treasury yields for 1-year, 2-year, and 10-year Treasuries in response to even-week dummies.

Table 2.8. Daily change in Treasury yield, 1994-2003

	1-year Treasuries	2-year Treasuries	10-year Treasuries
Dummy=1 in Week 0	0.000352	0.00208	-0.00167
	(0.13)	(0.61)	(-0.49)
Dummy=1 in Week 2,4,6	0.000208	-0.000445	-0.00191
	(0.09)	(-0.16)	(-0.71)
N	2608	2608	2608

t statistics in parentheses
* p<0.10, ** p<0.05, *** p<0.01

As can be seen, there is no meaningful response in any of the US Treasuries. One would expect that indications of accomodative monetary policy which result in excess stock returns would also, and perhaps more strongly, affect US Treasury yields. In Appendix B.1 Table B.1, I show that these results hold for the entire CMVJ sample (1994-2016).[14]

As an additional test, I explore whether my results for US Treasuries are consistent with the response of Fed Funds (FF) futures and Eurodollar (ED) futures. Again, one would expect futures contracts to react strongly to leaks about monetary policy. Table 2.9 below documents the daily change in the yields of the second FF futures contract, the fourth FF futures contract,

[14] CMVJ show similar results for the 2-year and 10-year Treasury. However, they combine the week 0 dummy with the week 2,4,6 dummy which therefore also captures the announcement effect and is a distinct mechanism from Fed leaks following biweekly board meetings.

and the fourth ED futures contract in response to even-week dummies from 1994 to 2003.[15,16] These specific futures contracts are widely used in the monetary policy shock literature (e.g., Gürkaynak et al. (2007) and Gertler and Karadi (2015)).[17]

Table 2.9. Daily change in futures yield, 1994-2003

	2nd FF Contract	4th FF Contract	4th ED contract
Dummy=1 in Week 0	-0.000251	-0.00104	0.00250
	(-0.18)	(-0.56)	(0.52)
Dummy=1 in Week 2,4,6	-0.00112	-0.000698	-0.00274
	(-0.94)	(-0.46)	(-0.74)
N	2607	2607	2607

t statistics in parentheses
* p<0.10, ** p<0.05, *** p<0.01

The results for futures are entirely consistent with those for US Treasuries as they simply do not respond in any statistically significant way to the even-week dummies. In Appendix B.1 Table B.2, I show that these results also hold for the entire CMVJ sample (1994-2016).

Together, Tables 2.8 and 2.9 cast doubt on the narrative in CMVJ that leaks following board meetings pushed up stock prices as one would expect these to also impact fixed income assets. However, these tables do not provide conclusive evidence against the CMVJ narrative. If the information leaked following board meetings is a commitment to accommodate *if needed*, then it would likely result in market expectations of a lower Fed Funds rate in bad states of the world. However, the commitment alone may make bad states less likely which may raise market expectations of the Fed Funds rate. The net effect could therefore be negligible. Indeed, CMVJ make precisely this argument, adding that these two effects of a commitment to accommodate

[15] I do not include the first FF futures contract as it is a current month contract and may expire prior to the subsequent FOMC meeting.

[16] While one might argue that prior to 2000, there is much less trading volume in FF futures and therefore to exclude that period, the pre-2000 period has been used extensively in estimating monetary policy shock series. For example, Gertler and Karadi (2015) use FF futures from as early as 1991.

[17] Gürkaynak et al. (2007) specifically argues that ED futures are the best predictor of fed funds rate at horizons beyond six months and as effective as predicting the fed funds rate as FF futures at horizons of less than six months.

and a lower probability of bad states would both positively impact stock prices. Therefore, given the evidence in this section casts some doubt on the CMVJ narrative but is not conclusive, I explore the potential for outliers in the next section.

2.4.3 Weekly Averaging and Outliers

The even-week effect is calculated based on weekly returns which average over the five business days in a week. This averaging process may result in a lot of important information being lost and may also mask potential outliers. Therefore, in this section, I examine the effect of FOMC cycle *days* rather than weeks and explore the impact of outliers over the period 1994-2003, where the even-week result may still hold.

First, I run a separate regression for each day of the cycle where I regress daily excess returns on a dummy for that *day* only. Figure 2.7 plots the coefficients of each of the regressions as well as their confidence intervals.[18] Each CMVJ-defined week is represented by a different colour. Figure 2.7 shows that only two days during even weeks are significant at the 5% level. Day 13 with a daily excess return of 26 bps and day 20 with a daily excess return of 45 bps. While a few more days are significant at the 10% level, these include the days around the announcement (which is likely related to the announcement effect rather than the even-week effect) and the days at the end of the cycle. Given board meetings are not on exactly the same day of the FOMC cycle throughout time, one might have expected that the individual days that make up even weeks would be consistently positive, even if not highly statistically significant. However, this does not appear to be the case in Figure 2.7, where a few outliers that occurred during even weeks clearly seem important. Therefore, this analysis suggests that even though the even-week result holds during the period 1994-2003, it may not necessarily be the result of board meetings and Fed leaks.

[18] I also document the regression results in a table. See Appendix B.1 Table B.4.

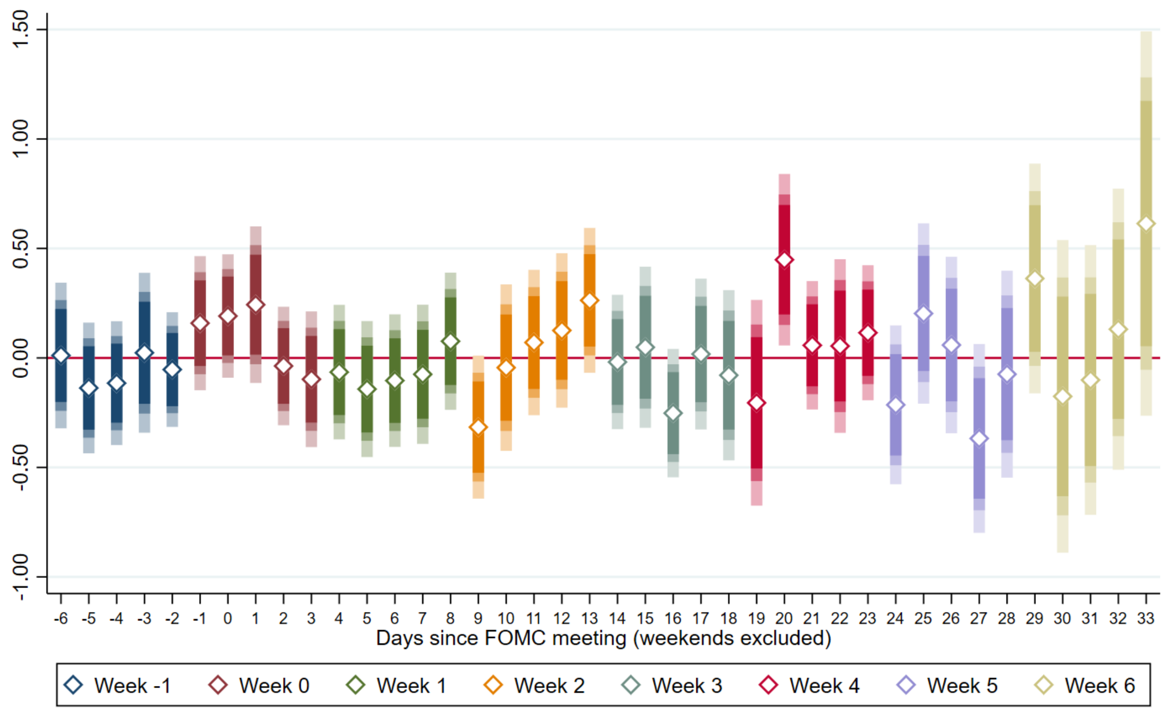

Figure 2.7. Regressing Daily Excess Return on Day of the Cycle, 1994-2003

Notes: The 40 regressions underpinning this chart have excess returns as the dependent variable. The only regressors are a constant and the dummy for the specific day of the cycle. The diamonds reflect the coefficient on the day dummy of each regression. The confidence intervals are shown by coloured bars. The darkest shade represents the 99% confidence interval, one shade lighter represents the 95% confidence interval, and the lightest colour represents the 90% confidence interval.

While Figure 2.7 clearly documents potential outliers, it does not provide evidence that these outliers are important for the even-week result nor that they are unrelated to board meetings. One possible, and indeed valid, counterargument is that these two days, while outliers, were outliers *because* of the biweekly board meetings. Therefore, in order to ascertain whether these outliers were indeed driven by board meetings, I examine financial market developments around day 13 and day 20 in greater detail.

I begin by noting that over the period 1994-2003, the average daily excess return is 0.03%, the average daily excess return for even weeks is 0.10%, the average daily excess return for day 13 is 0.28%, and the average daily excess return for day 20 is 0.47%. The single highest one-day return during this period occurred on day 20 of the FOMC cycle on 24 July 2002. On this day, the one-day excess return is 5.43%, over 50 times greater than that of the even-week average. The pertinent question is whether this one-day return was the result of a biweekly board meeting that leaked information about accomodative monetary policy.

First, I confirm that there was indeed a board meeting that took place prior to the excess return. Specifically, the meeting was on 22 July 2002, i.e., two days prior to the high return day. However, when examining financial market movements on that day, it appears the excess returns were *unrelated* to the board meetings. For context, as highlighted by the financial press at the time, 22 July 2002 saw the second biggest one-day gain for the Dow Jones industrial average on record. The big move in stocks appeared to be driven by positive progress in relation to the Enron scandal. In fact, in a conference call at the time, J.P. Morgan reassured investors that it would have no problems signing off its account. Moreover, the chief investment officer at Walnut Asset Management commented that "J.P. Morgan's comments this morning were also a good first step toward recovering credibility, you had the Senate and the House agreeing on their fraud policy, and you had the Adelphia executives dragged out in handcuffs."[19]

While the events surrounding 22 July 2002 do not appear to be related to the Fed or the board meeting, I also examine Fed communications and behaviour at the time. In the minutes of the FOMC's meeting on 13 August 2002, the FOMC itself noted that "market participants focused their attention on further revelations of corporate malfeasance" in reference to the Enron scandal.[20] In terms of specific actions, the FOMC maintained the policy rate at the same level at both this meeting and the following meeting on 24 September 2002. The Fed's behaviour and the context of financial market movements strongly suggest that the excess return on 22 July 2002 was not due to leaks following the board member meeting a few days prior.

Continuing to focus on day 20, its second largest value is 4.1% which is within the top 0.5% of excess returns during the period 1994-2003. This occurred on 28 October 1997, the day after a global stock market crash due to the Asian economic crisis of 1997. Figure 2.8 below

[19] https://www.cnn.com/2002/BUSINESS/asia/07/24/wallst.rebound/
[20] https://www.federalreserve.gov/fomc/minutes/20020813.htm

shows excess returns over the course of 1997.

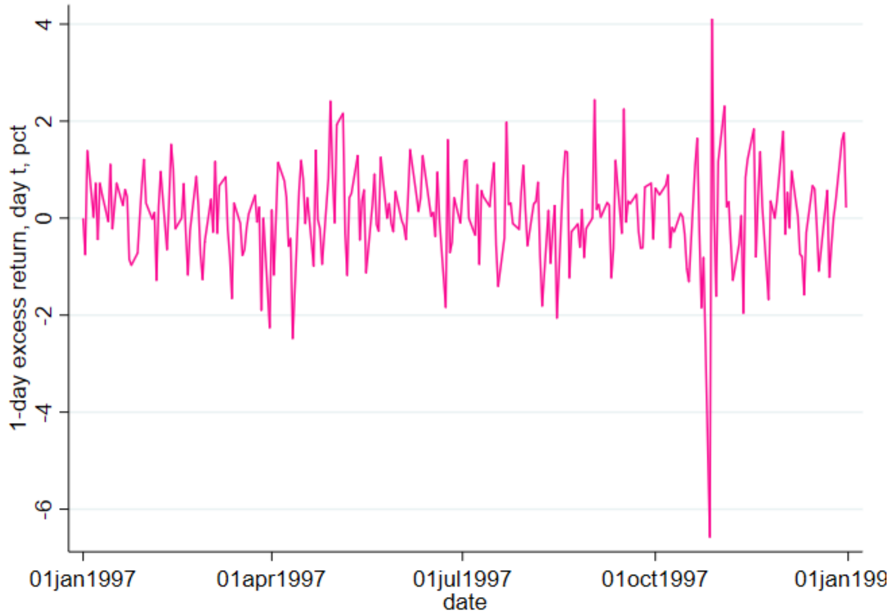

Figure 2.8. Excess Returns during 1997

It is clearly evident from the plot above that the stock market movements on 27 and 28 October 1997 are outliers. The crash on 27 October was related to the Asian economic crisis and, as the Securities and Exchange Commission noted in its Trading Analysis, the rally on 28 October was simply a rebound that was "consistent with the widespread view that the bounce-back in share prices reflected broad-based buying by institutional and retail investors."[21] Again, this does not appear to be related to specific Fed-related communication (informal or otherwise). Indeed, the FOMC noted in its minutes of the 12 November 1997 that while equity markets were volatile, short-term interest rates registered little change since the 30 September FOMC meeting.[22]

In a similar vein to my analysis of the day 20 outlier, I evaluate financial market developments specific to day 13. While the returns are not as large as those on day 20, the largest day 13 return is still within the top 1% of all returns. Specifically, the largest daily excess return

[21] https://www.sec.gov/news/studies/tradrep.htm
[22] https://www.federalreserve.gov/fomc/minutes/19971112.htm

is 3.74% and occurred on 11 October 2002. Moreover, this occurred four days after a board meeting but was still within an even week. Therefore, it would contribute to the even-week effect. As before, the relevant question for this paper is whether the return was due to leaks following the board meeting which would provide evidence supporting CMVJ. Looking at analyst reports and the financial press on this day suggests that the excess return was *unrelated* to the Fed but rather a response to unexpectedly positive earnings reports from many large firms.[23] The FOMC itself noted that "the subsequent release of better-than-expected news on profits for several major corporations buoyed equity prices" in the minutes of its 6 November 2002 FOMC meeting.[24] The second largest day 13 return is 3.21% on 7 January 2000. Similar to the largest day 13 return, this appeared to be in response to unexpectedly positive earnings reports, with analysts noting little movement in expectations of the Fed Funds rate.[25]

The narrative evidence documented above suggests that some of the most significant returns on day 13 and day 20 were unrelated to the Fed. However, it does not provide evidence that these two days were important for the overall even-week result. Therefore, to test whether these outliers actually drove the even-week effect, I re-estimate the even-week regression of column 1 Table 2.5 but simply exclude day 13 and day 20. The results are in Table 2.10 below.

Table 2.10. The Impact of Outlier Days, 1994-2003

	Full Sample	Excl. Day 13 & Day 20
Dummy=1 in Week 0	0.149**	0.149**
	(2.47)	(2.47)
Dummy=1 in Week 2,4,6	0.119**	0.0568
	(2.35)	(1.05)
N	2608	2460

t statistics in parentheses
* p<0.10, ** p<0.05, *** p<0.01

[23] See https://money.cnn.com/2002/10/11/markets/markets_newyork/index.htm and https://www.wsj.com/articles/SB1034334222879789356

[24] https://www.federalreserve.gov/fomc/minutes/20021106.htm

[25] See https://www.wsj.com/articles/SB947253922454914021 and https://rb.gy/p5655a

Table 2.10 has two columns. Column 1 restates the results without any exclusion (i.e., those of column 1 Table 2.5) for ease of comparison. Column 2 presents the results when day 13 and day 20 are excluded.[26] Once we exclude just two days, not only is the result no longer statistically significant, the coefficient has more than halved. Therefore, day 13 and day 20 appear necessary for CMVJ's main result. However, the largest returns on those days are inconsistent with CMVJ's proposed mechansim and appear to be statistical outliers.

As an additional robustness test, instead of removing day 13 and day 20, I winsorise the data. Specifically, I remove the top and bottom five percent of outliers. The purpose of this is to simply remove returns that are so large that they are likely the result of market events similar to those documented when examining day 13 and day 20. In Figure 2.9 below, I plot the distribution of the returns pre- and post-winsoring. While the means of the distribution are similar (0.29 vs 0.27), one can immediately see the long tails have been trimmed.

[26] Specifically, I drop day 13 and day 20 from the data and then run the same even-week regression. Note that this approach has the effect of reducing the sample size. Alternatively, I also redefine the even-week dummy for week 2 to include all days but day 13 and the even-week dummy for week 4 to include all days but day 20 to preserve the sample size. This alternative result produces a qualitatively similar result.

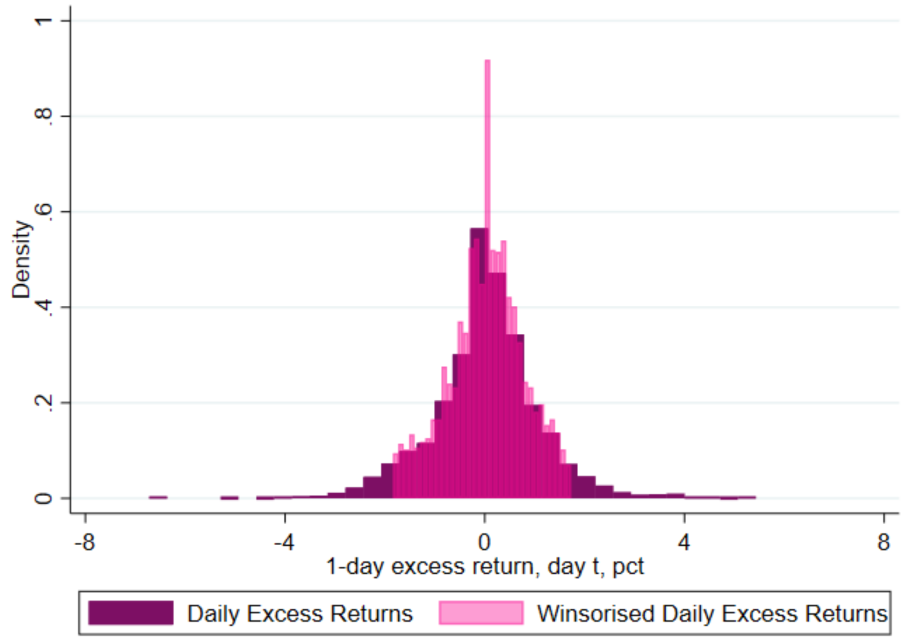

Figure 2.9. Histogram of Daily Excess Returns: Full Sample versus Winsorised

Next, in Table 2.11 below, I document the results of a regression with the winsorised stock returns compared to the full sample returns.

Table 2.11. The Impact of Outlier Excess Returns, 1994-2003

	Full Sample	Winsorised Sample
Dummy=1 in Week 0	0.149**	0.158***
	(2.47)	(3.58)
Dummy=1 in Week 2,4,6	0.119**	0.0567
	(2.35)	(1.64)
N	2608	2345

t statistics in parentheses
* p<0.10, ** p<0.05, *** p¡0.01

Similar to Table 2.10, Table 2.11 shows that the result disappears when removing outliers and the coefficient more than halves. When considering my finding in relation to outliers, alongside some of my earlier findings documenting the lack of robustness of both the even-week effect and the proposed mechanism from 2004 onwards, as well as the lack of an effect on fixed income assets, the totality of the evidence appears to suggest that the even-week result may have been a statistical artifact.

2.5 The Even-Week Result in an International Context

CMVJ show that international stock market returns also follow the FOMC cycle, though they do not test the robustness of this result in the same way as they do for US stock returns, nor do they explore central bank cycles of other countries. While there is work that explores stock returns around announcements of the domestic central bank for several countries (e.g., Brusa et al. (2020)), I am not aware of any work that explores the central bank cycle in other countries, as CMVJ have done for the US. Therefore, this section has two questions. First, is CMVJ's finding that international stock returns are driven by the FOMC cycle robust? Second, do central bank cycles, like the FOMC cycle, exist in other countries? In the case of the UK and Japan, I find that the answer to both questions is no.

2.5.1 BoE Cycle

Figure 2.10 provides preliminary answers to the two questions posed above. The chart on the left shows 5-day forward cumulative UK excess stock returns over the BoE cycle. There is no even-week result in the BoE cycle. However, it appears there is a type of pre-announcement drift.[27] The chart on the right shows UK stock returns over the FOMC cycle and appears to confirm CMVJ's finding that international stock returns are driven by the FOMC cycle.

[27] This would not be same as the pre-announcement drift documented by Lucca and Moench (2015) which would appear in week 0 when using CMVJ's week definition.

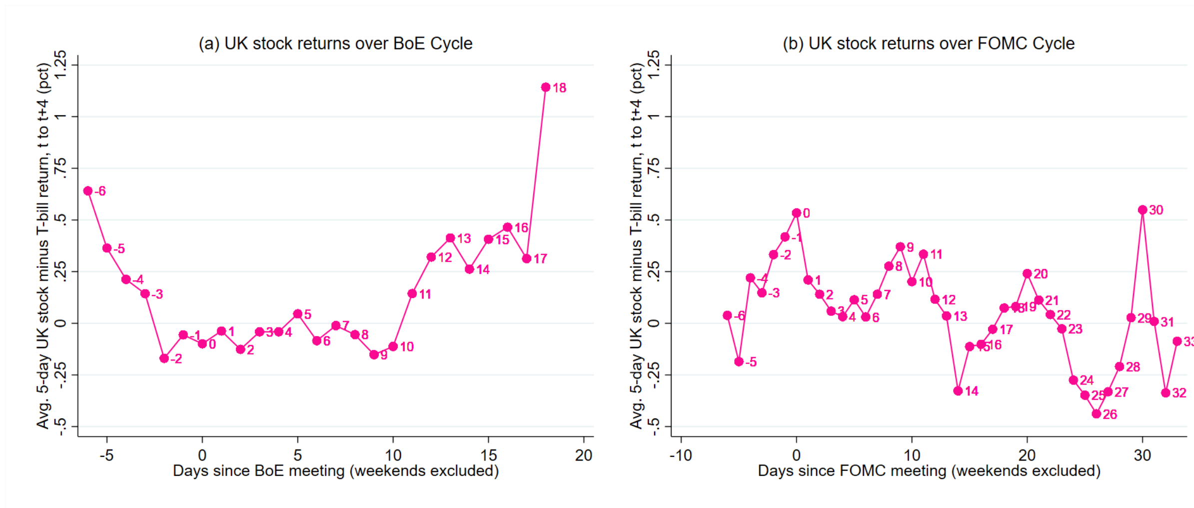

Figure 2.10. UK Stock Returns over the BoE and FOMC Cycles, 1997-2016

While the above visual evidence is useful, I run regressions to answer the two questions more rigorously. Table 2.12 shows the results of regressing daily UK excess stock returns on a combination of BoE and FOMC cycle dummies. I control for a type of pre-announcement drift given the evidence in the left chart of Figure 2.10. The FOMC week dummy is the standard CMVJ even-week dummy. Column 1 shows that only week -1 is significant (i.e., stocks rise in the week leading up to BoE meetings). Column 2 shows UK stock returns seem to spike in even weeks of the FOMC cycle, though the result is marginally significant. While this is consistent with CMVJ's findings, I find that after including the BoE week -1 dummy (i.e., controlling for the pre-announcement effect), the FOMC cycle is no longer significant in predicting UK stock returns (column 3). If CMVJ's hypothesis that the central bank cycle is a result of informal communication and leaks is correct, then this suggests that BoE is perhaps stricter about communication flows.

Table 2.12. Regressions of Daily Excess UK Stock Returns on BoE/FOMC Cycle Dummies

	(1)	(2)	(3)
Dummy=1 in BoE Week -1	0.143***		0.138***
	(3.74)		(3.59)
Dummy=1 in FOMC Week 0,2,4,6		0.0628*	0.0532
		(1.90)	(1.60)
N	4861	4861	4861

t statistics in parentheses
* p<0.10, ** p<0.05, *** p<0.01

2.5.2 BoJ Cycle

I repeat the above analysis, but this time in the case of Japan. Given the BoJ had much less consistency in the number of meetings per year, any interpretation should be treated with caution. The chart on the left of Figure 2.11 shows 5-day forward cumulative excess stock returns over the BoJ cycle. Like the BoE, there is no even-week result for the BoJ. The chart on the right shows Japanese stock returns over the FOMC cycle. While they do not follow the FOMC cycle as closely as UK stock returns, they appear to have even-week spikes.

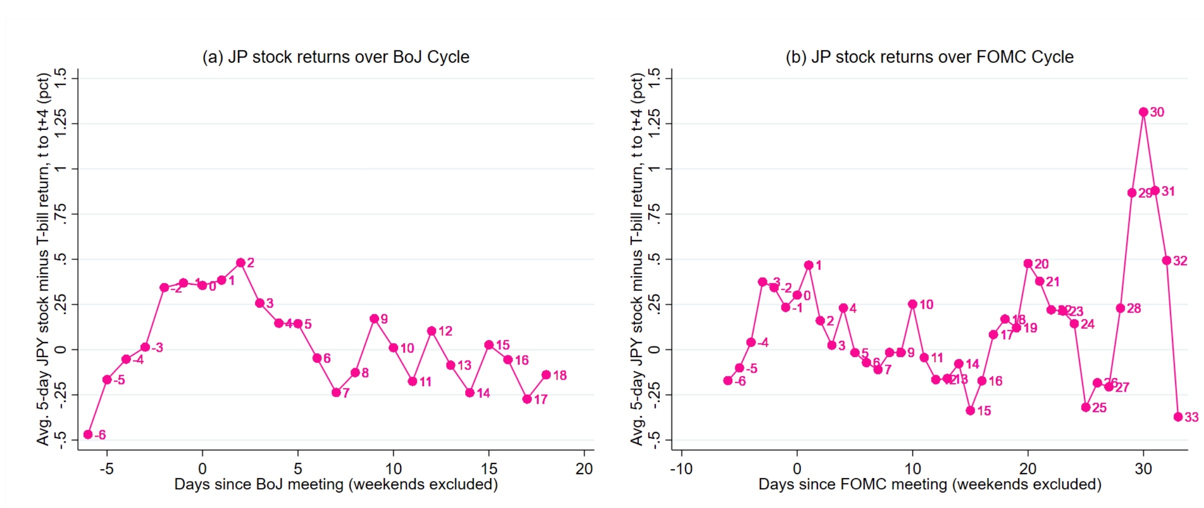

Figure 2.11. Japan Stock Returns over the BoJ and FOMC Cycles, 1998-2016

As before, I run regressions to answer the two questions more rigorously. Table 2.13 shows the results of regressing daily Japanese excess stock returns on a combination of BoJ and FOMC cycle dummies. I test for the pre-announcement effect (i.e., a dummy in BoJ week -1). The FOMC week dummy is the standard CMVJ even-week dummy. Column 1 shows that while week -1 is significant, it is negative. Column 2 shows that even without any controls, Japanese stock returns are not driven by the FOMC cycle. Again, if CMVJ's argument is correct, then the lack of a BoJ cycle suggests that like the BoE, and unlike the FOMC, the BoJ is potentially stricter about information flow.

Table 2.13. Regressions of Daily Excess JP Stock Returns on BoJ/FOMC Cycle Dummies

	(1)	(2)	(3)
Dummy=1 in BoJ Week -1	-0.156***		-0.156***
	(-3.45)		(-3.45)
Dummy=1 in FOMC Week 0,2,4,6		0.0188	0.0167
		(0.46)	(0.41)
N	4954	4954	4954

t statistics in parentheses
* p<0.10, ** p<0.05, *** p<0.01

2.6 Conclusion

CMVJ's even-week result, both in the US and internationally (i.e., that even weeks of the FOMC cycle predict excess stock returns), is very important in raising necessary questions about the conduct of central banks. Therefore, in this paper, I have evaluated some of CMVJ's key findings and make three key contributions. First, I find that the FOMC cycle does not drive US stock returns after 2004. Specifically, I find that the even-week result is only robust when one concludes the sample for analysis in 2016. When expanding the sample to end-2023, the result stops being robust as early as 2004. My second contribution evaluates the mechanism proposed by CMVJ which is that even-week returns are *causally* driven by leaks following biweekly board meetings at the Fed. I show that after 2004, board meetings were not biweekly, casting doubt on the underlying mechanism. Moreover, I show that before 2004, when the even-week result holds and board meetings are at least biweekly, there were important outliers that are unrelated to board meetings. When removing these outliers, the even-week effect no longer holds. My third contribution relates to evaluating the robustness of CMVJ's even-week result internationally. I find that the FOMC cycle does not drive UK excess stock returns when one controls for potential pre-announcement effects in the UK and that FOMC cycle does not drive Japanese stock returns regardless of whether one includes controls. I further show that neither the BoE or the BoJ, unlike the FOMC, have especially unusual patterns in their cycle.

Regardless of whether there is informal communication as posited by CMVJ, this area of research raises an important question: does informal communication (or targeted leaks) improve the effectiveness of monetary policy? While CMVJ's result may not be robust, they have raised the profile of this important question and documented creative ways to think about it.

2.7 Chapter Acknowledgements

Chapter 2 is currently being prepared for submission for publication of the material. The dissertation author was the sole investigator and author of this paper.

Chapter 3

Kites And Quails: Monetary Policy And Communication With Strategic Financial Markets*

3.1 Introduction

> *We might find ourselves having to act to stabilize the financial markets in order to stabilize the economy. None of us welcomes the charge that monetary policy contains a "Fed put," but, in extremis, there may be a need for such a put, if not in the strict sense of the finance term, then at least in regard to the direction of policy. How should we communicate about such actions, or, as used to be said of the lender of last resort, should we leave such actions shrouded in constructive ambiguity?* [1]

The Global Financial Crisis of 2007-8 challenged the consensus that monetary policy should focus primarily on inflation (Smets (2014)). Since then, there has been a significantly revived interest in whether monetary policy should account for financial stability concerns. Indeed, Woodford (2012) shows that loose monetary policy can increase financial instability and therefore central banks should include such concerns in their objective function.

In addition to such normative questions about central banks' objective functions, there

* Kites are a bird of prey and, like hawks, are in the Accipitridae family. Quail-doves are members of the Geotrygon bird genus in the pigeon and dove family.

[1] Governor Fischer, FOMC Meeting Transcript, March 15-16, 2016 (https://www.federalreserve.gov/monetarypolicy/files/FOMC20160316meeting.pdf)

are two sets of positive questions. First, do we see evidence from the revealed behaviour of central banks that they might already have financial stability as part of their objective function? Second, what are the implications of this type of objective function? There is substantial evidence that central banks account for the behaviour of financial markets when making decisions. We see this anecdotally from transcripts of policymakers' meetings, speeches, and other forms of communication. However, there is also more rigorous evidence of this. For example, Shapiro and Wilson (2022) find that "the FOMC's loss depends strongly on output growth and stock market performance and less so on their perception of current economic slack." Moreover, there is evidence that unexpected increases in interest rates lead to a significant decline in bank stock prices as well as a decline in short-term bank profits (e.g., Flannery and James (1984); Aharony et al. (1986); Alessandri and Nelson (2015); Busch and Memmel (2015); English et al. (2018); Altavilla et al. (2019); Ampudia and van den Heuvel (2019); Zimmermann (2019); Jiang et al. (2023); Uppal (2024).

Theoretically, there appear to be at least two common approaches in formalizing ways in which central banks care about financial stability. The first relates to *policy surprise*. For example, Stein and Sunderam (2018) argue that when central banks surprise the market, it can lead to costly volatility and subsequently instability for highly leveraged financial institutions. The second relates to a *directional preference* of the financial market participants whereby market participants might prefer low interest rates to high interest rates. This is often related to the 'Fed Put' (a term that describes a scenario in which the Federal Reserve keeps interest rates low in order to protect the stock market). Cieslak and Vissing-Jørgensen (2020) highlight that the Fed Put could result in moral hazard issues.

In our paper, we seek to understand the implications of these modified central bank objective functions when financial institutions are *strategic* players (i.e., they know the central bank cares about their stability). Using this game-theoretic set-up, we produce a number of contributions to the literature. First, we show that if investors (i.e., financial market participants) face costly readjustments when their investment position is not consistent with the central bank interest rate choice and the central bank internalizes part of this adjustment cost, then the central bank underreacts to shocks to the economy (see Appendix C.2 for a case study using Silicon Valley Bank and Signature Bank to illuminate this underreaction result). In a sense, this is consistent with the gradualism approach highlighted by Stein and Sunderam (2018). It is also similar to a result derived by Caballero and Simsek (2022). They show that even when the Fed does not agree with the market, because market expectations affect current asset valuations, it can induce the Fed to set an interest rate that partly reflects the market's view.

Next, we focus specifically on systemic financial institutions as most central banks have explicit mandates to ensure the stability of these institutions. Such institutions are often considered "too big to fail" and so arguably warrant additional supervision by regulators.[2] We derive a solution to the game where these systemic institutions have payoffs that relate to policy surprise and directional preference, and central banks are concerned about market readjustment costs. In this set-up, we show that when these institutions have market power, there is a policy distortion. Indeed, the designation of the institutions as being systemic relies on them having some measure of market power. Moreover, there is considerable empirical evidence documenting the growing market power of US banks (see, for instance Corbae and Levine (2022) and Corbae and Levine (2022)) and increasing evidence on the role of bank market power in the transmission of monetary policy (Drechsler et al. (2017)). Our market power-related distortion is akin to the moral hazard conjecture of Cieslak and Vissing-Jørgensen (2020). Put simply, given that the central bank is concerned with the stability of these institutions, they can make portfolio decisions that make it harder for the central bank to implement monetary policies misaligned with the interests of financial market participants.

We then explore two sets of policies to improve welfare. The first relates to classic papers on the appropriate degree of discretion for central banks. For example, Barro and Gordon (1983) shows that rules for central bank behaviour can be welfare-improving, but that given monetary policy is a repeated game between the policymaker and the private agents, it is possible

[2]See https://www.bis.org/bcbs/gsib/ for a list of designations.

that reputational forces might substitute for formal rules. Our findings relate more closely to the seminal paper by Rogoff (1985) which finds that social welfare can be improved if the appointed central banker is more hawkish than society (i.e., a central banker who differs from the social objective function as they place a larger weight on inflation-rate stabilization relative to employment stabilization). Akin to the hawk-dove result of Rogoff (1985), we derive a similar kite-quail result. Specifically, we can show that when the central bank can be of two types: quail (place greater weight on financial institution stability) and kite (place a greater weight on the real sector, i.e., the conventional central bank objective), it is socially optimal to have a central bank that is more kitish than the society.

Finally, we contribute to the growing literature on central bank communication. While there is already a large literature on central bank communication through policy announcements, the focus is largely on communicating a decision (Blinder et al. (2008)). Indeed, papers as early as Stein (1989) have explored the problems faced by the Fed in announcing its private information about future policies. Stein (1989) showed that the Fed can communicate some information about its goals through the cheap talk mechanism of Crawford and Sobel (1982). However, there is much less research on the effects of central bank communication *prior* to a decision being made. This is an especially important area given the prevalence of central bank speeches, press releases, and other announcements which are used to "guide" market expectations prior to the decision-making meeting as such communications can themselves influence the policy decision. While there is growing empirical evidence on both formal and informal communication (see, for instance Cieslak et al. (2019); Hansen et al. (2019); Bradley et al. (2020); Morse and Vissing-Jørgensen (2020)), as far as we are aware, only Vissing-Jorgensen (2020) explores the issue of this type of communication formally in a theoretical model. However, she does not allow for communication to have any benefit nor does she allow for the market to behave strategically, both features that seem *prima facie* important in describing communication between central banks and markets. We find that when communication is cheap talk, a maximally *kitish* central banker improves welfare more often than when the central bank can commit to full transparency. Interestingly, not only does the central bank better achieve its stability goals when a kitish central banker is appointed, but investors can also benefit, due to more transparent communication and market stability. This finding is novel and non-obvious, as the appointment of a *kite* has two effects on the expected payoffs of financial market participants, pushing in opposite directions. On the one hand, it discourages attempts to influence the central banks' policy, forcing market players to behave more as "policy takers," which reduces the conflict of interest with policymakers in the communication stage. As in Crawford and Sobel (1982), a smaller conflict of interest allows for more information transmission, making all parties – including market players – better off. This benefit of a kitish central banker is *in addition to* the aforementioned benefit of appointing a central banker more kitish than society, and therefore highlights a novel interaction between the central bank stance and the effectiveness of communication. On the other hand, holding the informativeness of communication fixed, financial institutions expect larger losses under a kitish central banker: the latter is less responsive to financial interests and more responsive to shocks that occur after communication, resulting in more surprises and instability. As a result, markets benefit from a central banker who puts little weight on their stability if the stability gain from transparent communication is large enough to compensate for the losses they incur due a lower policy influence.

3.2 Setting

We model the interaction between the central bank (CB) and a set of market players I using the following stylized three-period setting. At time $t = 0$, CB privately observes a shock $\omega_1 \sim U(-\phi_1, \phi_1)$ to the economy-stabilizing policy (i.e., the interest rate consistent with stabilizing the economy), and sends a public message $m \in M = [-\phi_1, \phi_1]$ potentially informative about ω_1. At $t = 1$, after having observed ω_1, each market player $i \in I$ simultaneously chooses an investment position $x_i \in X$, based on their expectation of the central bank interest rate, which is observed by the CB. Finally, at $t = 2$, CB observes a second shock to the economy stabilizing policy, $\omega_2 \sim U(-\phi_2, \phi_2)$, independent from ω_1, and chooses the policy rate $r \in \mathbb{R}$.[3] The economy-stabilizing policy at the time of the decision is $\omega \in [-\phi, \phi]$, where $\omega = \omega_1 + \omega_2$ and $\phi = \phi_1 + \phi_2$. After the policy decision is made, the game ends and payoffs are realized.

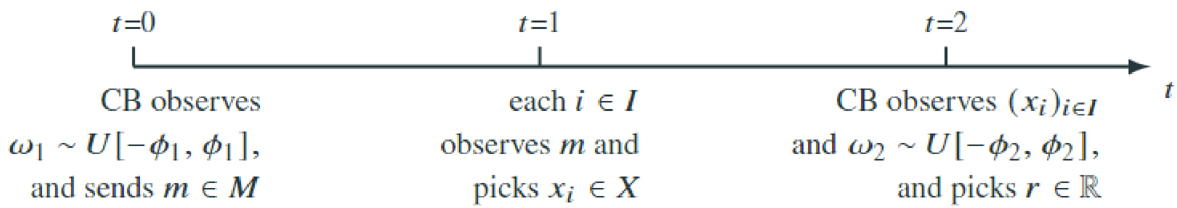

Figure 3.1. Timeline

Figure 3.1 provides a graphical representation of the timing. The interpretation of the timeline is as follows. The policy decisions of the Fed are made during meetings of the Fed Open Market Committee (FOMC), occurring once every six weeks, a period of time known as the "FOMC cycle." We can think of our simple game as occurring within the FOMC cycle, which is distinct from much of the existing literature that considers decisions at the time of announcement. In our game, $t = 2$ is the end of the cycle, when the policy decision is made and announced. Within the cycle, official communication occurs at fixed dates, and is meant to increase the transparency of the decision process of the FOMC, as well as provide future guidance to markets (see, e.g., Faust (2016) for an extensive discussion of the role of communication). With a simplifying assumption, we assume that policy-relevant communication occurs only once within the cycle, at $t = 0$, where at the time of communication, the policymaker has some uncertainty about the economy-stabilizing policy (captured by ω_2). This simplification is consistent with the fact that the speeches of the FOMC chair and vice-chair occur once per cycle, and that the committee observes a mandatory "blackout period" before the occurrence of the end-of-cycle meeting. Finally, in our timeline, investors make the relevant changes to their investment positions before the policy decision r is made. This modelling choice captures the idea that before the decision is made public, at $t = 1$ market players choose their investment positions

[3]The uniform distribution assumption is without loss of generality as far as ω_2 is concerned. Assuming that ω_1 is also uniformly distributed simplifies the communication analysis of section 3.4.1.

based on their future policy expectations. It allows for "forward guidance" and highlights the economic value that investors derive from correctly predicting future monetary policies.

3.2.1 Market Surprises and Underreaction

We make two main assumptions that shape the strategic incentives of investors and policymakers in our setting. First, the payoffs of the investors are lower when the mismatch between their investment positions and the policy decision r is greater. Second, the central bank internalizes at least part of this market adjustment cost.

Assumption 1 (Costly Market Readjustments) *Let $X = \mathbb{R}$ and x_i denote the element of X chosen by i. For $a, b \in \mathbb{R}$, let $S(a,b) = (a-b)^2$.[4] We assume that*

(i) the payoff function of investor $i \in I$ takes the form

$$u_i(x_i, r) = f^i(S(x_i, r), r)$$

for $f^i : \mathbb{R}^2 \to \mathbb{R}$ with $f_1^i < 0$, where f_1^i is the derivative of f^i with respect to its first argument;

(ii) the loss function of the central bank takes the following form. For $\alpha \in (0,1)$,

$$L_{CB}(x, r, \omega) = \frac{1}{2}(1-\alpha)S(r, \omega) + \frac{1}{2}\alpha g(\bar{S}(x, r)),$$

for $g : X^I \times \mathbb{R} \to \mathbb{R}$ with $g' > 0$, $g'' \geq 0$ and $\bar{S}(x,r)$ a weighted average over $\{S(x_i, r)\}_{i \in I}$.

The assumption has a straightforward interpretation. First, (i) means that the payoff of investor i is, *ceteris paribus*, decreasing in the mismatch $S(x_i, r)$ between her investment position x_i and the policy r. We can think of x_i as representing a choice from a continuum of portfolios of financial securities (e.g., stocks, bonds, and derivatives), where portfolio $x_i = r$ is the one performing

[4]We use the quadratic loss functional form for tractability as well as consistency with the form of central bank loss functions in the literature.

best at $t = 2$ if the policy is r. At $t = 1$, S is minimized in expectation when i holds a position x_i matching the expected future rate, and, when the investor picks a portfolio consistent with a given expected future rate, the further the realized rate is from this expectation, the worse their portfolio performs ex-post.[5] In this sense, S is akin to a *policy surprise*. We assume that this type of portfolio performance enters the payoff function of market players, picking up an incentive for markets to correctly predict future monetary policies.

Second, assumption 1(ii), when paired with $\alpha > 0$, means that the central bank internalizes the cost of the overall market readjustment due to a policy surprise. The interpretation is analogous to Stein and Sunderam (2018): unexpected policies can create large fluctuations in asset prices, undermining the financial health of highly leveraged institution.[6] As in Stein and Sunderam (2018) and Vissing-Jorgensen (2020), we interpret our loss function L_{CB} as the welfare loss associated to monetary policy, so that $\alpha > 0$ means that market instability due to surprises has a negative welfare effect over and beyond the traditional channels of monetary policy.[7] This assumption captures the fact that, since the Global Financial Crisis, central banks across the world have increasingly adopted explicit financial stability objectives which typically include monitoring the profitability of banks as well as the volatility of asset prices (Calvo et al. (2018)).[8][9] While this assumption reasonably reflects central bank mandates across a growing number of jurisdictions, there is also direct empirical evidence that central banks are concerned about stock market performance (see Shapiro and Wilson (2022) for evidence in the case of the

[5]We simply assume that this ex-post loss is increasing in the mismatch between their portfolio and the realized policy rate in a quadratic fashion.

[6]There is a large empirical literature showing that monetary policy shocks affect asset prices (e.g., see Rigobon and Sack (2004) for the effects on asset prices more generally and English et al. (2018) for the effects on bank stock prices specifically).

[7]This assumption is relaxed if we interpret L_{CB} simply as the loss of the CB, and reframe the welfare implications of the next sections as payoff-implications for the CB.

[8]One example from the Federal Reserve comes from its May 2023 Financial Stability Report. The Fed highlighted that higher interest rates could lead to credit losses for banks exposed to commercial real estate and that as a result it would expand its monitoring of banks with such exposures. See https://www.federalreserve.gov/publications/files/financial-stability-report-20230508.pdf

[9]The Bank of England specifically publishes the indicators it uses for macroprudential purposes. These include the return on assets for the banking sector and the VIX index (a measure of market expectations of 30-day volatility as conveyed by S&P 500 stock index options prices. See https://www.bankofengland.co.uk/-/media/boe/files/core-indicators/countercyclical-capital-buffer.xlsx for the full list of indicators.

Federal Reserve).[10] In the next section, we explore the case when these investors are systemically important financial institutions, where the argument that the central bank is concerned for their stability is even stronger.

Not surprisingly, our first result is that when the central bank cares about market surprises, it underreacts to shocks in the economy-stabilizing policy that occur after communication. By doing so, policymakers reduce market surprises and the associated welfare cost.[11] This simple intuition is captured by proposition 1, which follows immediately from the game structure and assumption 1(ii).

Proposition 1 (Underreaction) *Let σ be any strategy profile played in a PBE of the game, and let $r^\star(\sigma,\omega_1,\omega_2)$ be the corresponding on-path policy rate chosen by CB, expressed as a function of the shocks realizations (ω_1,ω_2). On the equilibrium path, the central bank under-reacts to the shock ω_2, that is $0 < \frac{\partial r^\star(\sigma,\omega_1,\omega_2)}{\partial \omega_2} < 1$.*

In the remainder of the paper, the focus of our analysis will be on a specific payoff specification that satisfies the requirements of assumption 1 and, we believe, delivers interesting insights and policy implications.

3.3 Interaction with Systemically Important Institutions

We start by analysing the case where the central bank cares about the financial stability of a set I of systemically important institutions. As highlighted earlier, many central banks, particularly those in advanced economies, have an explicit financial stability mandate on top of their price stability mandate. Indeed, accounting for the stability of systemic institutions is a core part of this mandate and is consistent with central banks conducting enhanced regulation and

[10]Although there are undoubtedly many sources of financial stability risks, we focus purely on risks to investors from unexpected changes in interest rates given the focus of this paper is on how including a financial stability objective affects the interest rate choice of the central bank.

[11]This feature of the policy decision is in line with the gradualism and underreaction that characterise equilibrium policies in Vissing-Jorgensen (2020) and Stein and Sunderam (2018), who also study monetary policy with costly policy innovations.

monitoring of these institutions due to the potentially systemic consequences of their instability. Moreover, central banks typically conduct stress tests of these systemic institutions to ensure that they are sufficiently stable, and in particular, that they are able to withstand adverse shocks (including unexpected changes in interest rates).

These systemic institutions have a structural interest rate exposure related to their business model that makes their profits sensitive to unexpected changes to the direction of the policy r.[12] In particular, we specify the following forms for the players' payoffs. For each $i \in I$ and $x_i, r \in \mathbb{R}$,

$$u_i(x_i, r) = \underbrace{-\frac{1}{2}(x_i - r)^2}_{\text{policy surprise}} \underbrace{-\beta r}_{\substack{\text{directional} \\ \text{preference}}} \tag{3.1}$$

where $\beta > 0$, implying that rate cuts positively impact investors' payoffs.[13] We interpret $\beta > 0$ as investors having a *structural* positive interest rate exposure deriving from their primary business, meaning that their profits are typically negatively impacted by rate hikes.[14]

The main policy implications presented in this section do not depend on the sign of β.[15] While the sign does not matter, we require that, on average, $\beta \neq 0$.[16] If β was zero, it would imply banks have no directional preference for interest rates. While Drechsler et al. (2021) show that in response to changes in interest rates, bank profits are relatively stable, when they specifically consider unexpected changes in interest rates, they find that bank stock prices

[12]We will sometimes refer to r as a policy change. This is similar to assuming that $r = 0$ at the beginning of the game.

[13]One can consider the exposure arising from maturity transformation as being determined by β while the exposure from all other investments as being determined by the portfolio choice x_i. Alternatively, one can consider β as capturing a bank's exposure to credit risk via unexpected changes in interest rates while x_i captures its direct exposure to interest rate risk. In both cases, we define β such that a positive value implies higher rates negatively impact banks.

[14]Begenau et al. (2021) show that the US banking sector has typically been characterised by a positive interest rate exposure over the last two decades and English et al. (2018) show that bank stock prices are negatively impacted by unexpected rate hikes.

[15]A key difference is that a $\beta > 0$ would imply that rates are declining over time while $\beta < 0$ would suggest that rates are rising over time.

[16]As documented by Begenau et al. (2021), there is heterogeneity in interest rate exposure across banks. Their results hold on average and, importantly, for large banks (market makers). In our model, we can indeed allow for heterogeneity across investors, in terms of size and sign of β. We believe that our main results would still hold in such a richer setting, provided that β is on average positive (or different than 0) for systemic investors.

actually fall. This stock price result is consistent with the findings of English et al. (2018). In addition to a fall in bank stock prices, English et al. (2018) show that unexpected rate hikes lead to a decline in bank profits.[17] Uppal (2024) finds a similar result and provides evidence that the underlying mechanism driving the overall decline in profits is due to higher credit losses from rising delinquencies.[18] This latter result is consistent with considering β as a bank's exposure to credit risk via unexpected changes to interest rates (see footnote 13). Therefore, despite the sign of β being inconsequential for our results, for the purposes of interpretation, we set $\beta > 0$. Indeed, a positive β is also consistent with the fact that banking crises are typically followed by reductions in the interest rate set by the central bank to, among other things, help support the banking sector.

The loss function of the central bank depends on its ability to stabilize the economy and market instability due to a policy surprise. Specifically, for $\omega \in [-\phi, \phi]$, $x \in \mathbb{R}^I$ and $r \in \mathbb{R}$,

$$L_{CB}(x, r, \omega) = \underbrace{\frac{1}{2}(1-\alpha)(r-\omega)^2}_{\text{conventional loss}} + \underbrace{\alpha \frac{1}{2N}\sum_{i \in I}(r - x_i)^2}_{\text{market readjustment}}. \tag{3.2}$$

where $N = |I|$. The loss function is comprised of two terms. The first term is akin to what appears in the standard objective function of New Keynesian models (e.g., Galí (2015)). Hence, if the welfare weight of systemic investors' stability was $\alpha = 0$, the policy $r = \omega$ would achieve the optimal balance of inflation and economic slack. Note that this standard term already includes the change in banks' profits due to their traditional-business exposure β, which could be considered as part of the bank lending channel of monetary policy. The second term in the central bank loss function captures the welfare cost of a financial market readjustment due to policies that surprise large investors. For simplicity, we assume that each systemically relevant institution carries equal

[17] Much of the literature finds that unexpected interest rate hikes lead to bank stock price declines and short-term bank profit declines (e.g., Flannery and James (1984); Aharony et al. (1986); Alessandri and Nelson (2015); Busch and Memmel (2015); Altavilla et al. (2019); Ampudia and van den Heuvel (2019); Jiang et al. (2023)

[18] Similarly, Zimmermann (2019) finds in a repeated cross-section of 17 countries and 145 years that unexpected interest rate hikes result in greater loan losses and lower credit growth for banks, resulting in a decline in bank profits.

weight in contributing to this instability (i.e., they have equal market power), so that the relevant welfare inefficiency is a simple average of the readjustment costs of the investors in I.

The solution concept that we adopt for the equilibrium analysis of this simple game is Perfect Bayesian Equilibrium (PBE).[19] We restrict our attention to equilibria in which the central bank choice r at $t = 2$ is not contingent on the announcement $m \in M$ at $t = 0$. Considering equilibria in which the central bank conditions the policy choice r on previous communication would not yield additional insights, in the absence of reputational considerations. Hence, a strategy for the central bank consists of a communication rule $\sigma_m : [-\phi_1, \phi_1] \to \Delta(M)$, mapping realizations of ω_1 to a distribution over messages in M, and a policy rule $\sigma_r : X \times [-\phi_1, \phi_1] \times [-\phi_2, \phi_2] \to \Delta(\mathbb{R})$, mapping investors' positions and realizations of (ω_1, ω_2) to distributions over policy decisions. A strategy for market player i is an investment plan $\sigma_{x_i} : M \to \Delta(\mathbb{R})$, mapping the central bank's messages to investment positions. We write $\sigma_x = (\sigma_{x_i})_{i \in I}$ and denote a generic strategy profile $(\sigma_m, \sigma_x, \sigma_r)$ by $\sigma \in \Sigma$.

Before moving to a general solution for this game, it is useful to identify the communication strategies, investment rule and policy rule that would maximize expected welfare, as this will provide a useful benchmark for equilibrium welfare analyses.

3.3.1 First Best: The Competitive Benchmark

What would socially optimal communication, investment and policy look like? Given our assumption that L_{CB} is equivalent to the relevant measure of welfare loss in the society, it suffices to ask how the central bank would communicate, invest and set monetary policy if, similarly to a social planner, she could directly choose the investment rule σ_x. The main result of this section is that the ex-ante optimal strategy profile is a PBE of the game where investors chose their strategies autonomously, provided that the banking sector is perfectly competitive.[20] Before

[19] We do not need to impose additional requirements to achieve subgame perfection, since in all proper subgames CB knows all past shocks realizations.

[20] In what follows, we interpret N as competitiveness of the banking sector. For instance, we can think that the central bank cares about the largest institutions that cumulatively represent some fixed share of the total market size, so that N is the number of such institutions.

stating this result formally and discussing it, it is useful to introduce two additional pieces of notation. First, for each $\sigma \in \Sigma$, let $W(\sigma) = -\mathbb{E}[L_{CB}|\sigma]$ denote the ex-ante welfare when the strategy profile is σ. Second, we let $\hat{W}$ denote the highest achievable welfare in the game if the central bank (or a social planner) could select any strategy profile from Σ, that is

$$\hat{W} = \max_{\sigma \in \Sigma} W(\sigma) \tag{3.3}$$

and let $\hat{\Sigma}$ be the argmax of (3.3).

The following proposition describes how such "first" best strategy profiles look like, and it identifies the case where a profile in $\hat{\Sigma}$ arises in a market equilibrium.

Proposition 2 (Competitive Benchmark) *$\hat{\Sigma}$ is the set of all strategy profiles presenting each of the following three properties:*

(i) Communication is fully informative, namely σ_m partitions $[-\phi_1, \phi_1]$ in singletons: different messages are sent in different states.

(ii) Investment is unbiased, namely $\sigma_{x_i}(m) = \mathbb{E}[\omega_1|m, \sigma_m]$ for each $i \in I$ and each $m \in M$ played with positive probability in equilibrium;

(iii) The central bank balances its objectives on the path of play. For $\bar{x} \equiv \frac{1}{N}\sum_{i\in I} x_i$, and every $(\omega_1, \omega_2) \in [-\phi_1, \phi_1] \times [-\phi_2, \phi_2]$,

$$\sigma_r(\omega_1, \omega_2, x) = (1-\alpha)\omega + \alpha\bar{x},$$

holds if $x_i = \omega_1$ for all $i \in I$.

Additionally, there exist an equilibrium profile σ^{com} such that $(i), (ii)$ and (iii) hold at the limit for $N \to \infty$.

The proof is in the appendix. The three properties of proposition 2 have intuitive interpretations: the central bank communicates all available information ω_1 about its' expected future policy, investment positions minimize expected readjustments given the expectations about the economy-stabilizing rate ω which are formed after communication, and the final policy is chosen to optimally balance stabilization of the economy and of financial markets. Perhaps more subtly, (i), (ii) and (iii) taken together imply that in this type of game, $\hat{W}$ does not depend on N, so that looking at the mismatch between $\hat{W}$ and the ex-ante expected welfare achievable in equilibrium for different level of competitiveness N is meaningful.

The main implication of the proposition is that when there is perfect competition between systemic investors (i.e., $N \to \infty$), there is an equilibrium σ^{com} that achieves the maximum welfare $\hat{W}$ that could be obtained by a benevolent planner. The intuition is as follows. First, it is easily verified that in every PBE strategy profile (including σ^{com}), the central bank uses the same policy rule $\sigma_r^\star$, that selects a rate lying between ω and $\bar{x}$, optimally balancing the two stability objectives. In particular, $\sigma_r^{com}(\omega_1, \omega_2, x) = \sigma_r^\star(\omega_1, \omega_2, x) = (1-\alpha)\omega + \alpha\bar{x}$, which satisfies property (iii) of proposition 2. Second, in a perfectly competitive market, investors *individually* behave as policy takers. While they anticipate that policymakers will take into account the *aggregate* market position $\bar{x}$ in setting the rate, each investor is too small to influence aggregate outcomes. This implies $\bar{x} = \mathbb{E}[\omega|m, \sigma_r^{com}]$. To see why, note that if the average position of the market was different from the expected real economy-stabilizing policy, that is $\bar{x} \neq \mathbb{E}[\omega|m, \sigma_r^{com}]$, then investor i would expect the central bank to follow $\sigma_r^\star$ and choose an intermediate policy, so that $\mathbb{E}[\sigma_r^{com}(\omega_1, \omega_2, x)|m, \sigma_r^{com}, x] \neq \bar{x}$. However, as policy takers, they would also find it individually rational to change their position to match perfectly this expected r, minimizing the surprise loss. Hence, $\sigma_{x_i}^{com}(m) = \mathbb{E}[\omega|m, \sigma_r^{com}]$ must hold in equilibrium, satisfying (ii). Finally, note that this "unbiased" investment plan is optimal also from the point of view of policymakers, who use communication to guide investors to select $x_i = \omega_1$, the investment position that makes it easier to stabilize the economy without disrupting markets in the future. Full guidance to $x_i = \omega_1$ is only achieved with fully transparent communication, by choosing $\sigma_m^{com}(\omega_1) = \omega_1$ for all $\omega_1 \in [-\phi_1, \phi_1]$, or any other communication rule that satisfies (i).[21]

[21]Even when investors behave as policy takers, there are multiple partition equilibria with partial information

Policy and Payoffs in the Competitive Benchmark

The on-path policy rate for any competitive strategy profile σ^{com}, and in general for any $\sigma \in \hat{\Sigma}$, is

$$r^{\star}(\sigma, \omega_1, \omega_2) = \omega_1 + \underbrace{(1-\alpha)\omega_2}_{\text{underreaction}}. \tag{3.4}$$

Transparent communication implies that the central bank fully reacts to policy target shocks occurring before communication. In contrast, and consistently with proposition 1, the welfare-optimal policy exhibits underreaction to state innovations that were not previously communicated, to limit market surprises.

Simple algebra, finally, leads to the following expressions for ex-ante welfare and investors' expected payoffs, for each $\sigma \in \hat{\Sigma}$, including the competitive limit case σ^{com}:

$$W^{com} = W(\sigma) = -\frac{1}{2}\alpha(1-\alpha)\sigma_2^2 \tag{3.5}$$

$$EU_i^{com} = \mathbb{E}[u_i|\sigma] = -\frac{1}{2}(1-\alpha)^2\sigma_2^2 \tag{3.6}$$

where $\sigma_2^2 = var(\omega_2)$. Note that both expressions are decreasing in σ_2, proportional to the residual policy uncertainty after communication and the resulting market instability. Additionally, the first-best average welfare is the lowest when the welfare weights of real-sector instability and financial market instability are close ($\alpha \approx \frac{1}{2}$), while, intuitively, investors are ex-ante better-off if the central bank puts high weight on market stability.[22]

The results of this sections seem optimistic, as the first-best average welfare can be obtained assuming competitive markets. This assumption, unfortunately, is likely to be violated in reality, where systemically important institutions – almost by definition – have large size and market power and there is considerable empirical evidence documenting market power in the financial sector (see e.g., Drechsler et al. (2017) and Corbae and Levine (2022)). To see what

transmission. Standard arguments, however, lead to the selection of the fully transparent equilibrium.

[22]We are not modeling here the long-term financial market losses due to an unstable economy. In this sense, we might think of our investors as "short-term" oriented.

happens as the banking sector becomes less competitive, we explore the general solution of the game for $N < \infty$.

3.4 Oligopolistic Competition

When firms have market power they are typically able to influence market prices and quantities through their individual decisions. In our setting, while we did not model the economy explicitly, large systemic institutions have a similar influence over equilibrium monetary policy. To see this, note that in all PBE the central bank plays the optimal policy rule $\sigma_r^\star(\omega_1, \omega_2, x) = (1-\alpha)\omega + \alpha\bar{x}$. As shown in the appendix, such a rule provides the unique minimizer of L_{CB} in any possible subgame reached at $t = 2$. Fixing this rule, for each $i \in I$, $(\omega_1, \omega_2) \in [-\phi_1, \phi_1] \times [-\phi_2, \phi_2]$ and $x \in X^I$ it indeed holds

$$\frac{\partial \sigma_r^\star(\omega_1, \omega_2, x)}{\partial x_i} = \frac{\alpha}{N} > 0$$

which is non-negligible when I is finite. Intuitively, when investors have market power, their *individual* losses due to financial assets that perform badly under a candidate policy rate matter for aggregate market stability and therefore for the policy decision.[23] Hence a change in the individual position x_i affects the choice of r. We refer to $\frac{\alpha}{N}$ as the *policy influence* of an individual investor.

It is worth first briefly considering whether this policy influence is plausible and importantly whether it is more likely for systemic investors (those with weight α in the central bank loss function) relative to non-financial corporations (those with weight $1-\alpha$ in the central bank loss function). While it is difficult to measure policy influence directly, one can certainly consider whether central banks engage more with financial institutions than non-financial institutions as this would suggest there is at least a forum for possible influence. First, central banks engage

[23]This argument is analogous to that in the granularity literature where idiosyncratic firm-level shocks can explain an important part of aggregate movements given the firm size distribution (see e.g., Gabaix (2011)).

considerably more with systemic institutions given their regulatory mandate. Specifically, these institutions receive enhanced monitoring and are subject to addition regulations (e.g., the systemic risk buffer). Second, analysing entries in the calendars of Federal Reserve governors between 2007 and 2018, Morse and Vissing-Jørgensen (2020) document a striking finding: governors met with financial institutions and financial interest groups more than three times as much as with non-financial institutions and non-financial interest groups (1,573 interactions versus 507 interactions). Together, these facts provide some suggestive evidence that financial institutions have relatively more engagement with central banks and given that central banks explicitly care about their stability, financial institutions may as a result have greater policy influence.

To better understand how investors' direct policy influence changes the strategic behaviour of investor i, let us focus on the simple case in which the central bank has revealed ω_1 at $t=0$ and $x_j = \omega_1$ for each $j \neq i$. The expected utility of investor i at $t=1$ given $\sigma_r^\star$ and $\bar{x}_{-i} = \omega_1$, is

$$\mathbb{E}[u_i(x_i,r)|\omega_1,\bar{x}_{-i},\sigma_r^\star] = \underbrace{-\frac{1}{2}(1-\alpha)^2\sigma_2^2 - \beta\omega_1}_{\substack{\text{expected payoff} \\ \text{from unbiased investment}}} \underbrace{-\frac{1}{2}\left(1-\frac{\alpha}{N}\right)^2(x_i-\omega_1)^2}_{\substack{\text{expected readjustment due} \\ \text{to investment bias}}} - \beta \underbrace{\left[\frac{\alpha}{N}(x_i-\omega_1)\right]}_{\substack{\text{policy distortion due} \\ \text{to investment bias}}}. \tag{3.7}$$

We immediately notice two differences relative to the perfectly competitive case ($\frac{\alpha}{N} \approx 0$). First, the expected readjustment due to investment bias is smaller, since the policy influence of i makes the central bank choose a policy rate closer to x_i than if the investor had no market power. Second, and relatedly, investment bias leads to a policy distortion that matters on top of its effect on the performance of portfolio x_i, because of the positive interest rate exposure $\beta > 0$.

This second effect is key for the next result, which shows how biased investment can be used strategically in equilibrium to distort policies in the desired direction.

Proposition 3 (Oligopoly Distortions) *The oligopolistic equilibrium investment plan σ_x^{oli} en-*

tails investment bias. More generally, fix any communication rule σ_m and let σ_x^{oli} be such that each investor best responds to $(\sigma_m, \sigma_r^\star)$ and to all other investors strategies. Then,

(i) Investment decisions are consistent with a lower interest rate relative to the one expected to stabilize the economy. The bias magnitude is strictly increasing in α and β and strictly decreasing in N. For each $i \in I$ and $m \in M$

$$\sigma_{x_i}^{oli}(m) - \mathbb{E}[\omega_1 | m, \sigma_m] = -\frac{\alpha\beta}{(1-\alpha)(N-\alpha)}.$$

(ii) If the central bank fully reveals ω_1, that is, $\sigma_m = \sigma_m^{com}$, than the on-path policy decision exhibits a downward distortion. For each $\omega_1 \in [-\phi_1, \phi_1]$ and $\omega_2 \in [-\phi_2, \phi_2]$,

$$r^\star(\sigma^{oli,tr}, \omega_1, \omega_1) = r^\star(\sigma^{com}, \omega_1, \omega_2) - \frac{\alpha^2\beta}{(1-\alpha)(N-\alpha)}$$

where $\sigma^{oli,tr} = (\sigma_m^{com}, \sigma_x^{oli}, \sigma_r^\star)$, with σ_m^{com} fully informative about ω_1.

To see why investment bias arises in equilibrium, consider again the case where the central bank always fully reveals ω_1 and in which all investors other than i have chosen unbiased positions $x_j = \omega_1$. The first order condition in equation 3.7 is

$$\underbrace{\left(1 - \frac{\alpha}{N}\right)^2 (\omega_1 - x_i)}_{\substack{\text{marginal cost from} \\ \text{downward bias}}} = \underbrace{\frac{\alpha\beta}{N}}_{\substack{\text{marginal benefit from} \\ \text{downward bias}}}$$

which yields the bias of proposition 3 if $N = 1$.[24] The marginal benefit from a downward investment bias is positive for I finite, and therefore $x_i = \omega_1$ is not optimal for an individual investor, even when all other institutions are choosing unbiased investment. And this is more true the higher the individual policy influence $\frac{\alpha}{N}$ and the higher the benefit from a rate cut β. It is easy

[24] Recall that for the sake of developing intuition we assumed $x_{-i} = \omega_1$. This will not be true in equilibrium, so that the expression of the equilibrium bias for a generic N, reported in proposition 3, is different from the one we obtain from this special case.

to see that in equilibrium, each systemic institution will choose portfolio positions consistent with smaller rate hikes than those really expected, potentially increasing their own losses from a restrictive change in monetary policy. This result, apparently surprising, is consistent with the evidence by Begenau et al. (2021), who find that the 4 largest US banks typically hold net derivative positions reinforcing (instead of hedging) their interest rate exposure.

Why would investors take risky positions that create even more losses in case of adverse rate shocks? The key is that, by behaving as if the rates were to remain lower than optimal for the economy, the systemic institutions in our game make it harder for the central bank to make large and undesired policy changes. This is the result shown in the second part of proposition 3 if we assume commitment to fully informative communication: the on-path rate decision is systematically lower than the one that maximizes welfare. As such, our results provide one possible microfoundation for the empirical evidence on interest rate exposure.

Moreover, our results are similar in spirit to some of the findings in the literature around risk-taking of 'too-big-to-fail' institutions (which we refer to as systemic institutions in this paper). For example, Afonso et al. (2014) show empirically that banks deemed too-big-to-fail engage in greater risk-taking because they believe they will be rescued if they fail. However, taking greater risk, all else equal, increases the likelihood of their failure and so pushes the regulatory authority to provide more protective measures (e.g., asset guarantee programs) than it would have otherwise. Similarly, in our paper, systemic institutions increase their exposure to losses from higher rates which pushes the central bank to be more accomodative by cutting rates or raising rates less than it would have otherwise. A recent example that highlights the pressure the Federal Reserve faces in relation to raising interest rates was during the banking panic of 2023. During this time, the former chair of the Federal Deposit Insurance Corporation (one of the institutions with a financial stability mandate in the US) publicly called on the Federal Reserve to stop raising interest rates.[25] Indeed, during this banking panic, both the Fed's behaviour and market expectations were consistent with the underreaction prediction of our model (see

[25] See https://www.cnn.com/2023/03/13/investing/sheila-bair-svb-fed-rates/index.html

Appendix C.2 for details). As we will see in the next section, a similar result holds, *ex-ante*, in the strategic communication equilibrium.

A consequence of these distortions is that when the central bank commits to full information transmission at $t = 0$ and players best respond, the society incurs an ex-ante welfare loss relative to perfect competition. In particular, it holds

$$W(\sigma^{oli,tr}) - W(\sigma^{com}) = -\frac{\alpha^3\beta^2}{2(N-\alpha)^2(1-\alpha)} < 0.$$

Investors benefit from their policy influence, achieving a higher ex-ante payoff than in the perfectly competitive case. Indeed, it can be shown that, for each $i \in I$,

$$EU_i^{oli,tr} - EU_i^{com} = \frac{(2N-1-\alpha)\alpha^2\beta^2}{2(N-\alpha)^2(1-\alpha)} > 0.$$

As one would expect, both these gaps close if the degree of competition N increases, if the welfare weight α of market stability decreases, or if the structural exposure of systemic investors β increases.

We have shown that when the market is not perfectly competitive, investment is generally biased, and, in the case when the central bank communicates all available information about the economy-stabilizing policy, there is, on average, a welfare loss for the society. Can more communication flexibility close this gap relative to perfect competition? Or will it instead lead to further welfare losses? We address this question in the next section.

3.4.1 Communication

We now relax any assumption about commitment to transparent communication and look at how equilibrium strategic communication will look like in our simple model, and what its welfare implications will be. An unfortunate implication of proposition 2 and proposition 3 is that communication cannot bring us back to the ex-ante optimum $W(\sigma^{com})$. In fact, in any oligopolistic PBE, investment is biased downwards relative to $\mathbb{E}[\omega_1|m,\sigma_m^{oli}]$, the market expectations of the future economy-stabilizing rate given any equilibrium communication rule σ_m^{oli} and the announcement m made at $t=0$. This means that the unbiasedness requirement of proposition 2 is violated so that $\sigma^{oli} \notin \hat{\Sigma}$. Even more importantly, given the cheap talk nature of comunication without commitment in this environment, flexibility ends up creating an additional welfare loss relative to commitment to transparent communication.

It is easy to show that equilibrium communication has the exact same characteristics as the communication strategy of the seminal cheap talk game studied by Crawford and Sobel (1982), where the receiver sender bias $b = \frac{\alpha\beta}{2\phi_1(N-\alpha)(1-\alpha)}$.

Proposition 4 (Cheap Talk) *Any communication rule played in an oligopolistic PBE partitions $[-\phi_1,\phi_1]$ in a finite number of intervals, so that equilibrium communication is never fully informative. In addition:*

(i) A PBE strategy profile $\sigma_P^{oli} = (\sigma_{m,P}^{oli}, \sigma_x^{oli}, \sigma_r^{\star})$ corresponding to P partition elements, produces the following ex-ante welfare:

$$W(\sigma_P^{oli}) = W(\sigma^{oli,tr}) - \frac{1}{2}\alpha(1-\alpha)\hat{\sigma}_{1,P}^2$$

where $\hat{\sigma}_{1,P}^2$ is the residual variance of ω_1 induced by the equilibrium communication strategy.

(ii) The residual variance $\hat{\sigma}_{1,P}^2$ is decreasing in the absolute value of the investment bias, and in the number of partition elements P. The maximum number of partition elements $\bar{P}$ is the smallest integer greater than or equal to

$$\frac{1}{2}\left(\sqrt{1+\frac{4\phi_1(N-\alpha)(1-\alpha)}{\alpha\beta}}-1\right)$$

which is non-decreasing in the absolute value of the investment bias.

As shown in the appendix, proposition 4 follows immediately from Crawford and Sobel (1982) and from the game structure once we impose individual rationality of investment and policy decisions.

The main implication of the result is that fully informative communication about policy intentions would be optimal for the central bank in this context, since welfare is decreasing in any the residual uncertainty about ω_1, as described in point (i). However, unfortunately, a fully informative equilibrium is not achievable. To see this, imagine that the *CB* was playing $\sigma_m(\omega_1) = \omega_1$ in a PBE, which is fully informative. Investors' would best respond with an investment plan $\sigma_x^{oli}(m) = m - \frac{\alpha\beta}{(N-\alpha)(1-\alpha)}$, to bias the future rate downwards. But the central bank would then be tempted to systematically announce higher rate hikes at $t = 0$, deviating to $\sigma_m(\omega_1) = \omega_1 + \frac{\alpha\beta}{(N-\alpha)(1-\alpha)}$. This credibility loss arises from the fact that the interests of the central bank and investors are not fully aligned and, when investors have policy influence, they do not behave as policy takers.

The consequence of the loss in information transmission relative to commitment to informative communication is that, on average, the central bank will provide less guidance, markets will be less ready for changes in policies, and policies will create more instability. This negative effect is minimized in the most informative equilibrium, the one with $P = \bar{P}$, since more partitions imply more informative communication. Not surprisingly, point (ii) of proposition 4 shows that the informativeness of communication in this society-preferred equilibrium increases the closer the investment distortion is to zero. Hence, the maximum amount of guidance achieved in equilibrium by the central bank increases in the competitiveness of the banking sector and decreases in the welfare weight of market surprises and in exposures of investors.

Note that investors too are ex-ante worse-off when the central bank does not communicate clearly. When the equilibrium strategy profile is σ_P^{oli}, the systemic investors' loss relative to commitment to full information is

$$EU_i^{oli,tr} - EU_i^{oli,P} = \frac{1}{2}(1-\alpha)^2 \hat{\sigma}_{1,P}^2 > 0 \tag{3.8}$$

in fact, being less able to predict future policies, they will be find it harder to balance the need to follow the central bank and the one to influence it.

Finally, notice that, for any equilibrium with $P \leq \bar{P}$ the on-path policy decision distortion will be negative on average,

$$\mathbb{E}[r^{\star}(\sigma_P^{oli}, \omega_1, \omega_2) - r^{\star}(\sigma^{com}, \omega_1, \omega_2)] = -\frac{\alpha^2 \beta}{(N-\alpha)(1-\alpha)}$$

but positive for some realizations of ω_1 (e.g., those very close to the upper bound of the respective partition element). This is due to the partial information transmission of the cheap talk equilibrium.

One of the key insights of this section is that the central bank faces a trade-off between avoiding surprises *ex-post* and communicating effectively: on the one hand, market stability concerns ($\alpha > 0$) push policymakers to systematically underreact to state innovations ω_2 that could surprise markets. Such ex-post policy adjustments can be desirable for the reasons discussed in our first-best analysis: market surprises generate welfare-detrimental instability. On the other hand, the greater the stability concerns α, the larger the investment bias and the communication loss: the central bank's ex-post effort to avoid policy surprises reduces the effectiveness of forward guidance, a force that increases instability.

3.4.2 Towards a Kitish Central Bank

We have shown that systemically important investors can use market power to influence the policies of the central bank, and that this creates an inefficiency and a welfare loss, especially when the central bank has private information that cannot be communicated credibly. What remedies could we use to fix these issues?

In the beginning of our discussion of the oligopoly, we have introduced the concept of policy influence, which we measured by $\frac{\alpha}{N}$ the direct effect of an individual investor behaviour on the central bank choices. A positive policy influence is what makes the oligopoly different from perfect competition, and, indeed, what drives the inefficiency due to investment and policy distortions. An obvious way to do so would be reducing the scale of systemically important institutions. This would likely be a long and difficult process, that might create inefficiencies due to economies of scale, not modeled here, and indeed as mentioned earlier, concentration appears to be increasing over time. The second possibility would be to reduce α. While we cannot reduce the *welfare weight* of a market disruption, we can consider what would happen if we were to appoint a (representative) central banker with an objective function different from societal welfare. In particular, in a spirit similar to Rogoff (1985), we conduct the following thought experiment: imagine that we could appoint a central banker who has an objective with the same functional form as in (3.2), but with a relative weight $\tilde{\alpha} \in [0,1]$ of market stability potentially different from α. What would the optimal central banker look like?

Formally, let us denote by $\sigma^{oli,tr}(\tilde{\alpha})$ and $\sigma^{oli}_{\bar{P}}(\tilde{\alpha})$, respectively, the full communication oligopolistic strategy profile and the cheap talk oligopolistic equilibrium when the weight in the loss function of the central bank is $\tilde{\alpha}$. We then ask what is the level of $\tilde{\alpha}$ that maximizes $W(\sigma^{oli,tr}(\tilde{\alpha}))$ or $W(\sigma^{oli}_{\bar{P}}(\tilde{\alpha}))$, where the function W, as in the previous section, maps strategy profiles into ex-ante welfare levels (computed based on the welfare weight α).

To better illustrate the answer, let us denote a central banker with $\tilde{\alpha} < \alpha$ as *kitish* and a central banker with $\tilde{\alpha} > \alpha$ as *quailish*. Finally, we denote a central banker with preferences equal to the society ($\tilde{\alpha} = \alpha$), as unbiased. A kitish central banker focuses on the real economy, paying relatively litle attention to market positions, while a quailish central banker is more accommodating towards financial markets. Our main result is that it would be optimal for the society to appoint a kitish central banker.

Proposition 5 (Kites) *The optimal central banker is kitish. In particular, the following holds:*

(i) There exist thresholds $\sigma_2^{oli}, \sigma_2^{oli,tr} > 0$ on the non-communicable uncertainty such that: under transparency, maximally kitish central bankers ($\tilde{\alpha} = 0$) improve welfare over unbiased ones if and only if $\sigma_2 < \sigma_2^{oli,tr}$; under cheap talk, they improve over unbiased ones if and

only if $\sigma_2 < \sigma_2^{oli}$. Moreover, it has $\sigma_2^{oli} > \sigma_2^{oli,tr}$.

(ii) The optimal central bankers are such that $\tilde{\alpha}^{oli}, \tilde{\alpha}^{oli,tr} \in (0, \alpha)$, where $\tilde{\alpha}^{oli,tr}$ is the central banker under transparency and $\tilde{\alpha}^{oli}$ is the central banker under cheap talk. The welfare loss under transparency and $\tilde{\alpha}^{oli,tr}$ is smaller than the loss under cheap talk and $\tilde{\alpha}^{oli}$. Additionally, $\tilde{\alpha}^{oli,tr}$ strictly positive and increasing in α, N, and ϕ_2.

(iii) When the central bank commits to transparency, the market ex-ante payoff is always strictly larger when the central banker is unbiased than when she is kitish. However, if communication is cheap talk, then the market ex-ante payoff can be larger under a kitish central banker than under an unbiased one.

The intuition for why society might prefer a kitish central banker is simple. As discussed above, a kitish central banker reduces systemically relevant investors' influence over policies, and is more welfare-enhancing the larger the market power of investors. If the choice is between an unbiased central banker and one that assigns no weight on market stability, part (i) tells us that the latter will be desirable when state innovations occurring after communication are of limited magnitude (σ_2^2 low), especially when a kite can also improve information transmission (i.e., under cheap talk). If $\tilde{\alpha}$ can be fine-tuned, the optimal central banker is always kitish, but not maximally so: the reason why the optimal weight is non-zero is that market surprises are costly for society ($\alpha > 0$), so that some degree of market stabilization is always optimal.

What is perhaps more surprising is the difference highlighted by part (iii) of the proposition, that is, under cheap talk a kite can even improve the performance of financial institutions. The key intuition is that when misalignment of incentives compromises information transmission (in the cheap talk equilibrium), a more kitish central banker improves communication: a kitish central banker reduces the marginal net benefits from investment downward bias, making it optimal for the investors to choose a portforlio position x_i more aligned with their true expectations about ω. Consistently with proposition 4.ii, this bias reduction means that more information is communicated in equilibrium, and the resulting reduction in policy uncertainty has a stabilizing effect on financial markets, making *both* the society and the financial institutions better off. If this gain in information transmission is large enough, market players are willing to give up some of their policy influence to achieve it (as is the case when a kitish central banker is appointed). The following example provides a stark illustration of this phenomenon.

Babbling Monopoly

Let $N = 1$, and $\phi_1 = \frac{1}{2}$. It is easy to verify that in this monopolistic setting $\frac{\alpha\beta}{(1-\alpha)^2} > \frac{1}{4}$ implies that $\bar{P} = 1$, so that $\hat{\sigma}^2_{1,\bar{P}} = \frac{1}{12}$ in the cheap talk equilibrium. Under an unbiased central banker, the monopolistic investor's ex-ante value of the game is $EU^{\alpha} = -\frac{1}{2}(1-\alpha)^2\left[\frac{1}{12} + \sigma_2^2 + \frac{(\alpha\beta)^2}{(1-\alpha)^4}\right] + \left(\frac{\alpha\beta}{1-\alpha}\right)^2$, which simplifies to $-\frac{1}{2}(1-\alpha)^2\left[\frac{1}{12} + \sigma_2^2 - \frac{(\alpha\beta)^2}{(1-\alpha)^4}\right]$. With a maximally kitish central banker, i.e., $\tilde{\alpha} = 0$, the investors' ex-ante value of the game is $EU^0 = -\frac{1}{2}\sigma_2^2$ in the fully informative equilibrium. If $\frac{1}{12} - \frac{(\alpha\beta)^2}{(1-\alpha)^4} > \frac{1-(1-\alpha)^2}{(1-\alpha)^2}\sigma_2^2$ then $EU^0 > EU^{\alpha}$, and our monopolistic investor would give up all its policy influence to improve communication. Note that the condition is satisfied for $\frac{\alpha\beta}{(1-\alpha)^2} \in (\frac{1}{4}, \frac{1}{\sqrt{12}})$ and σ_2^2 small enough.

3.4.3 Game Repetition

In this section we return to the assumption that $\tilde{\alpha} = \alpha$, and we consider an infinite-horizon repetition of the game outlined in the previous section. Let $\tau = 0, 1, \ldots$ index the stage game repetition, and let $\omega^{\tau}, m^{\tau}, x^{\tau}, r^{\tau}$ denote the respective quantities in the stage game τ.

We address the following question: can an infinitely repeated interaction allow the central banker to better discipline markets, communicate more transparently and achieve the first best? We decide to focus on equilibria such that the investment decision x_i^{τ} is contingent only on past messages $(m^0, \ldots, m^{\tau})$ and investment decisions $(x_i^0, \ldots, x_i^{\tau-1})_{i \in I}$, and is therefore independent from central bank past policies, or shock realizations. [26]

[26]The set of equilibria of the repeated game depends on the observability of the shocks ω_1 and ω_2 at the end of the stage game. In particular, when the central bank equilibrium policies are contingent on shock realization, any market strategy treating histories where the central bank has deviated differently from histories where the central bank has played according to the equilibrium requires market players to draw inferences about actual shock realizations to detect such deviations. The relation between m^{τ} and r^{τ} can sometimes convey information about a central bank deviation, we decide to abstract from these situations to simplify the analysis and avoid imposing

We will see in the next propositions that, even within the restricted class of equilibria, repeated interaction can indeed achieve the first best on the path of play with a very simple type of PBE strategy profile, provided that forward guidance is valued enough by markets (i.e., ϕ_1 is high). At the same time, repeated interaction can facilitate collusion between large investors, sometimes with perverse consequences on welfare.

Proposition 6 (CB Disciplines Markets) *Consider the infinitely repeated game with discount factor $\delta \in (0,1)$. If $\phi_1 > \frac{\alpha\beta}{(N-\alpha)(1-\alpha)}\sqrt{3\frac{2N-\alpha-1}{1-\alpha}}$ then there exist a threshold $\delta^\star \in (0,1)$ such that if $\delta \geq \delta^\star$ the ex-ante stage welfare is at its first-best level on the equilibrium path of some PBEs. Such PBE strategy profiles are qualitatively equivalent to the following simple profile.*

(i) At $\tau = 0$, CB selects $m^\tau = \omega_1^\tau$. At $\tau > 0$, CB selects $m^\tau = \omega_1^\tau$ if $x_i^s = m^s$ for each $s < \tau$ and each $i \in I$, while CB draws $m^\tau \sim U(-\phi_1, \phi_1)$ after any other history.

(ii) At $\tau = 0$, each $i \in I$ selects $x_i^\tau = m^\tau$. At $\tau > 0$ each $i \in I$ selects $x_i^\tau = m^\tau$ if $x_j^s = m^s$ for each $s < \tau$ and each $j \in I$, while each $i \in I$ selects $x_i^\tau = -\frac{\alpha\beta}{(N-\alpha)(1-\alpha)}$ after any other history.

(iii) At each $\tau = 0, 1, \ldots$, CB sets rate $r^\tau = (1-\alpha)\omega^\tau + \alpha\bar{x}^\tau$.

The proposition has a very simple interpretation. If the variance of the shock ω_1 is large enough (ϕ_1 high), the value of forward guidance is high for market players. The central bank can leverage its informational advantage by conditioning fully informative guidance on market discipline – threatening to revert to meaningless communication if the directives are not followed in previous periods. The threat is credible, because when investors expect communication to be uninformative and stop listening, no deviation from the central bank can restore confidence in its announcements. In this equilibrium, policymakers play an active role in disciplining large banks. After investors' deviations, the central bank does not simply revert to the most efficient stage PBE communication rule presented section 3.4.1: market discipline follows from a threat

additional assumptions.

to revert to maximally inefficient communication. This threat allows for market discipline, even in cases where investors would be better off under repetition of the stage game equilibrium $\sigma_{\bar{P}}^{oli}$.

It seems natural, at this point, to ask whether the first best can also be sustained by spontaneous coordination of market players, with policymakers playing a less active role. To address this question, we restrict the attention to equilibria that satisfy two additional requirements. First, we impose that the central bank communicates as transparently as possible given the expected equilibrium investment bias $\mathbb{E}[\bar{x}^{\tau} - \omega_1^{\tau}|x^0,...,x^{\tau-1},m^0,...,m^{\tau-1}]$ at the history reached. The focus on most informative equilibrium communication in every contingency can be seen as the dynamic equivalent of what we did in section 3.4.1 in a static setting. Second, we focus on equilibria where the ex-ante expected stage payoff of each investor $i \in I$ is weakly greater than the expected payoff of the static game equilibrium $\sigma_{\bar{P}}^{oli}$. We call these equilibria *collusive*, as they represent a (weak) Pareto improvement for large investors relative to the static setting (where there is no coordination), and the central bank plays a passive role in disciplining markets.

The next proposition suggests that under certain conditions there exist equilibria where the oligopolists coordinate on the competitive equilibrium, but repeated interaction can also lead to detrimental effects on stability of the financial markets and the real economy.

For each $n = 1, ..., N$, and fixed parameters $\alpha, \beta, \phi_1, \phi_2$, let $\hat{\sigma}_{1,\bar{P}}^2(n)$ denote the residual variance of ω_1 after communication in the most informative stage-game equilibrium when the market consists of n large investors.

Proposition 7 (Market Collusion) *Consider the infinitely repeated game with discount factor $\delta \in (0,1)$. The following holds:*

(i) If $\hat{\sigma}_{1,\bar{P}}^2(N) > \left[\frac{\alpha\beta}{(N-\alpha)(1-\alpha)}\right]^2 \frac{2N-\alpha-1}{1-\alpha}$, then there exist $\delta_1^{\star} \in (0,1)$ such that, if $\delta \geq \delta_1^{\star}$, the first-best ex-ante welfare is the ex-ante stage welfare of some collusive PBE of the infinitely repeated game.

(ii) If $\hat{\sigma}_{1,\bar{P}}^2(1) - \hat{\sigma}_{1,\bar{P}}^2(N) < \left[\frac{\alpha\beta}{(1-\alpha)^2}\right]^2 - \left[\frac{\alpha\beta}{(N-\alpha)(1-\alpha)}\right]^2 \frac{2N-\alpha-1}{1-\alpha}$, then there exist $\delta_2^{\star} \in (0,1)$ such that, if $\delta \geq \delta_2^{\star}$, the ex-ante welfare of the monopolistic static game is the ex-ante stage

welfare of some collusive PBE of the infinitely repeated game.

Investors are ex-ante better-off in the monopolistic equilibrium (ii) relative to the first best (i) if $\hat{\sigma}^2_{1,\bar{P}}(1) < \left[\frac{\alpha\beta}{(1-\alpha)^2}\right]^2$*, while they are better off at the first best if* $\hat{\sigma}^2_{1,\bar{P}}(1) > \left[\frac{\alpha\beta}{(1-\alpha)^2}\right]^2$.

The first part of the proposition suggests that if the policy uncertainty that can be resolved by forward guidance is sufficiently large – hence payoff-relevant for investors – then market players will benefit from self-discipline: investors refrain from using their market power to bias future policies, which allows the central bank to communicate all its information on future policy changes, achieving the first best.

The second part of the proposition serves as a caveat: unsurprisingly, collusion between investors might well push in the opposite direction, facilitating a stronger exercise of market power by strategic market players, allowing them to act collectively as a single large player.

It is possible that large institutions might try to coordinate on investment strategy that maximize their aggregate expected payoffs. The type of collusion that is optimal for large investors need not be efficient from the central bank perspective, nor need it be as inefficient as the monopolistic one. However, we believe that the two equilibria highlighted in proposition 7 are important focal points from the point of view of the analyst and, potentially, the policymaker. The end of proposition 7 provides a sufficient condition for the monopolistic collusive equilibrium to benefit large investors more than the first best. If the gains from the exercise of monopolistic market power are large enough, markets should not be expected to self-coordinate on the first best, even when the first best is among the feasible collusive PBE.

All in all, the (partial) analysis of the repeated game suggests that, dynamic incentives could lead to both higher or lower central bank losses. To maximize the chances of achieving efficient outcomes, policymakers might have to play an active role, threatening to withhold future guidance if large market players refuse to cooperate.[27]

[27] Extending the analysis to broader set of equilibria, including those where market strategies are contingent on previous monetary policy, would likely lead to further interesting implications. Large investors – and not only policymakers – could strategically exploit communication to further bias policy or increase the informativeness of

3.5 Discussion

We have shown that when the central bank cares about market stability and systemic investors, like large banks, are asymmetrically affected by rate hikes and by rate cuts, the latter can use their market power to make adverse rate changes less likely. Specifically, they can choose market positions that would create a costly readjustment if the central bank was to choose a high economy-stabilizing rate (i.e., a high policy rate). Intuitively, this policy distortion depends on investors' policy influence $\frac{\alpha}{N}$ which is decreasing in the degree of competitiveness of the market N, and increasing in the welfare weight of market instability α. The resulting welfare loss is even larger if the central bank cannot guide markets credibly, which is reasonable when the misalignment between markets and policymakers make the former diffident towards announcements. Increasing competition in financial markets (i.e., transitioning towards a larger number of smaller and "less systemic" institutions) and appointing a central banker who puts little weight on market stability might increase both information transmission and welfare in equilibrium. Repeated interaction is not guaranteed to resolve the conflict of interest: the central bank can try to discipline markets under the threat of withholding future guidance if the strategic investors exercise their policy influence. But, when forward guidance is not valuable enough, market discipline might not be feasible, and large institutions could instead increase their effective market power through oligopolistic collusion.

The simple analysis yields a number of predictions. It suggests that a central bank that cares about large investors is expected to tailor policies to the interest rate exposures of large investors, with this tailoring increasing in the concentration of the financial market. As a consequence, net interest rate exposures of the largest systemic investors should systematically predict the variation in interest rate choices not explained by economic fundamentals. Moreover, the model would also predict that the growing market power over time should result in greater

announcements. For instance they could threaten the central bank to ignore future forward guidance if previous announcement are not precise enough or previous rate changes are not accommodating enough.

underreaction by the central bank and, all else equal, a declining trend in the interest rate. Indeed, while not a causal claim, we do see that in the data rising bank market power since the 1980s has coincided with a steadily declining interest rate. While the effectiveness and informativeness of forward guidance are in general expected to be limited when systemic investors have large market power, market discipline and transparent communication should be more common when the uncertainty regarding economic fundamentals is sufficiently high.

3.6 Chapter Acknowledgements

Chapter 3 is currently being prepared for submission for publication of the material. Giampaolo Bonomi and the dissertation author, Ali Uppal, are co-authors of this chapter.

Chapter 4

The Propagation Of Microeconomic Shocks Through Financial And Production Networks

4.1 Introduction

There is a large literature documenting the importance of fully understanding the impact of credit supply shocks. The Global Financial Crisis emphasised this importance, but it also highlighted another crucial component for how such shocks are transmitted: interconnectedness. Firms are connected to each other through supplier-customer relationships and they are connected to banks through their need for external financing. Therefore, it should be no surprise that such interconnections can result in potentially substantial amplification of bank-level credit supply shocks. In this paper, we show that accounting for the structure of both financial and production networks is essential to understand the channels through which financial shocks transmit to the real economy.

Our paper makes two contributions. First, we develop a general equilibrium model featuring production networks, firm-bank relationships, and financial frictions. We tractably introduce a financial sector with heterogeneous banks into a benchmark model of production networks. While the model is static, the need for banks arises as production takes place before households consume, and a financial friction prevents firms from making any payments to households once production took place. In this environment, banks act as intermediaries providing loans to firms to pay for using the household's factors of production and financing these through deposits from households. Banks have market power and choose deposit and loan rates to maximise profits. The model features endogenous firm-bank relationships and bank-varying loan and deposit rates. Prices of goods produced by non-financial firms reflect interlinkages through the production network and the cost of financing determined by the competition between banks. The model delivers a framework that speaks to the observed features of the firm-bank network and captures new mechanisms through which idiosyncratic shocks to bank lending or firm productivity permeate through the financial and the production sector.

In the model, exogenous shocks to a banks' credit supply are captured through exogenous shocks to a bank's ability to finance itself. The corresponding partial equilibrium effect on firms' effective financing cost corresponds precisely to our proposed shift-share instrument. The aggregate effects on output, in turn, depend in part on the production network but also on the degree to which firms can switch between banks. Intuitively, the more substitutable banks are, the smaller is the impact of idiosyncratic bank-level shocks on firms' effective financing rate. We document that general equilibrium propagation depends on structural parameters governing competition in the banking sector as well as the leverage of banks. These features amplify the direct effects of credit supply shocks and play an important role theoretically in understanding how such shocks ultimately affect GDP.

An important insight of our model is that the aggregate effects of bank-level shocks can be characterised by a set of sufficient statistics that can be measured or estimated in the data. These include the structure of the bank-firm network, the Domar weights, the share of intermediate inputs in production, the share of bank profits in GDP, bank equity returns, bank substitutability in loan and deposit markets, market shares of banks, and a bank's leverage ratio. This characterisation makes our framework highly amenable to empirical exploration.

Our second contribution is to describe in detail how one could implement an empirical strategy guided by our theoretical framework. While we do not explicitly implement the empirical strategy, we detail precisely the data needed, the structural interpretation, and the empirical specification. Our suggested empirical setting exploits sizeable misconduct provisions by UK banks from 2012 to 2015 as plausibly exogenous shocks to bank capital and lending. Over this time, UK banks set aside around £87 billion in misconduct provisions, a sum that accounts for a sizeable share of aggregate capital and reflects the largest financial services redress exercise in UK history. The vast majority of the incurred provisions relate to the misselling of Payment Protection Insurance during the late 1990s and early 2000s. The resulting timing-difference in misconduct taking place and provisions being made suggests that misconduct provisions incurred in the 2010s are plausibly orthogonal to unobserved shocks affecting banks' lending behaviour (Tracey and Sowerbutts (2023)). We argue that variation in both timing and magnitude of misconduct provisions across firms can be used to estimate bank-specific credit supply shocks.

To construct firm-level credit supply shocks, we explain that one can use the constructed bank-level credit supply shocks with detailed information on bank-firm loan relationships from the Burea van Dijk FAME database. The FAME database provides annual accounts data for over a million registered UK firms and includes data on firm-bank relationships at the loan level. Given firms differ in terms of the number of banks that provide them with funding and the number of loans each bank provides, one can leverage differences in the strength of bank-firm relationships to construct, for each firm, measures capturing a firm's weighted exposure to the constructed bank-specific credit supply shocks. The weights are firm-specific and reflect the historical importance of a given bank as a supplier of funding for the firm. We argue that the resulting firm-specific shocks provide an instrument for a firm's credit supply.

The firm-level shocks fall into the class of shift-share instruments (Bartik (1991)). We frame our discussion of the underlying identification assumptions by showing that the instrument can be microfounded and structurally interpreted in the context of recent models studying firm-bank relationships in the presence of imperfect bank substitutability (Herreño (2023), Balloch and Koby (2023)). Specifically, we show the equivalence between our instruments and the structural responses of firm-level funding costs to bank-level shocks in our model. This is helpful to frame the assumptions required for us to correctly infer firm-level credit supply shocks in terms of structural primitives. Finally, we show how one can use the instruments in an empirical strategy following Di Giovanni et al. (2018).

Our paper contributes to four main literatures: the literatures on shock propagation in real and financial networks, the large literature studying the real effects on macroprudential policy, and the interaction between competition and financial stability. The fast growing literature on shock propagation has primarily evolved by separately studying the role of real and financial networks. For real networks, the focus has been on understanding the nature of supply chain propagation, while in the second case the focus has been mainly on banks. A key contribution of this paper is to bridge these two literatures by explicitly accounting for the role and determinants of bank-firm linkages, and how these affect shock propagation while also highlighting the role banking sector competition plays in this propagation.

The literature on the real effects of production networks in macroeconomics is large and mostly theoretical. Seminal work in recent years include Acemoglu et al. (2012) and Gabaix (2011) who provide asymptotic results showing when idiosyncratic shocks might pose aggregate risks, highlighting network asymmetry in the sectoral production network and granularity of firms respectively. Baqaee and Farhi (2019) analyse non-linear propagation effects in efficient, non Cobb-Douglas economies, and propagation in inefficient economies is addressed in, e.g., Baqaee (2018) and Bigio and La'o (2020). While quantifying propagation in empirical cross-sectional settings is challenging, a recent strand of the literature has used natural disasters in the US (Barrot and Sauvagnat (2016)) and Japan (Boehm et al. (2019), Carvalho et al. (2021)) coupled with detailed information on firm linkages to provide quasi-experimental evidence. We contribute to the literature on production networks by studying their relationship with firm-bank networks, theoretically and also describe a possible empirical application.

The literature on financial networks has primarily focused on the trade-off between risk-sharing and contagion, and it has built on the contributions by Allen and Gale (2000) and Freixas et al. (2000). Our focus on the bank-lending channel is more closely related to theoretical work (e.g., Holmstrom and Tirole (1997), Gertler and Kiyotaki (2010)), who highlight that credit supply shocks may lead to significant real effects. Cross-sectional evidence on the bank-lending channel is provided by, e.g., Khwaja and Mian (2008), Jiménez et al. (2012), Chodorow-Reich (2014), Jiménez et al. (2017), Amiti and Weinstein (2018), and Jiménez et al. (2020). This paper contributes to this strand of the literature by providing a theoretical analysis of the bank-lending channel accounting for production and financial networks as well as a proposed empirical strategy.

Two related contemporaneous papers analyse the impact of credit supply shocks on direct and indirect propagation using administrative data from Spain. Alfaro et al. (2021) focus on propagation through industry input-output relationships, while Huremović et al. (2023) focus on firm-to-firm propagation. Both papers identify firm-specific credit supply shocks through statistical decomposition of credit growth following Amiti and Weinstein (2018), and both use a framework featuring production networks and a distortion based on Bigio and La'o (2020) to analyse production network propagation.

Our paper differs in three important ways: First, we provide a theoretical framework that explicitly accounts for the role of banks, for how firms and banks form relationships, and for the joint determination of financial and real prices. In this sense, we are most similar to recent work by Heipertz (2021). However, unlike Heipertz (2021), we focus on the propagation of credit supply shocks. Second, given we explicitly model the banking sector, we are also able to contribute to the literature on the debate between financial stability and banking sector competition (e.g., Corbae and Levine (2024), Benchimol and Bozou (2024)). Indeed, Proposition 8 shows exactly how the propagation of shocks depends on structural parameters that determine competition in the banking sector. By explicitly modelling the banking sector, we hope to contribute to the literature by highlighting the importance of banks in the propagation of shocks. Third, guided by our theoretical framework, we propose an empirical strategy with instruments that have a clear economic interpretation as they can be traced back to unexpected shocks to bank equity (which can be thought of as isomorphic to macroprudential policy changes). Our theoretical model provides a structural interpretation of the proposed instrument, further aiding its economic interpretation.

The remainder of this paper is as follows. Section 4.2 provides our theoretical framework, including formalising the aggregate impact of granular bank lending shocks (Proposition 8). Section 4.3 describes a possible empirical application, including detailing two candidate credit supply shocks. Section 4.4 concludes. Appendix D.1 provides a proof of Proposition 8.

4.2 Theoretical Framework

In this section, we develop a simple general equilibrium model accounting for both firm-bank relationships and input-output networks. We consider a one-period model to highlight the mechanisms.

4.2.1 Setup

The economy is populated by three types of agents: a continuum of households $h \in [0,1]$; there is a large number of non-financial firms $i \in N$; and a large number of banks $b \in B$. While the model is static, events take place across two stages.

Timing

Production takes place at time 0. Each household is endowed with 1 unit of the factor of production, which we call labour. Labour is sold in a competitive market at price w at time 0. Firms require financing to purchase factors of production, which is provided by banks through loans. Banks fund these loans through deposits provided by households. Banks determine loan and deposit rates and firms select banks for financing.

At time 1, production finishes. Products are sold and consumed. Firms repay loans to banks, households withdraw their deposits, and markets clear.

Households

In period 0, a continuum of households $h \in [0,1]$ each inelastically supplies one unit of labour at the competitive wage w. Each household has logarithmic preferences over goods $i \in N$:

$$\mathscr{C}(c_1, \ldots, c_N) = \sum_{i=1}^{N} \beta_i \log\left(\frac{c_i}{\beta_i}\right) \tag{4.1}$$

where β_i measures expenditure shares across the available varieties and are normalised such that $\sum_i \beta_i = 1$. Given consumption does not take place until time 1 and households do not discount the future, households deposit all the cash they earned from supplying their labour at time 0. While households receive a return on their deposits at the deposit rate $R_{bd} > 1$, they also have idiosyncratic preference shocks over different banks. Therefore, the overall utility of household h from consuming $\mathscr{C}$ and having banked with bank b is

$$U_h(\mathscr{C}, b) = \mathscr{C} \varepsilon_h(b) \tag{4.2}$$

where $\varepsilon_h(b)$ is an idiosyncratic preference shock which captures household-specific reasons for a household to prefer one bank to another such as location, online banking services, and so on. The vector of shocks ε_h is drawn from a Fréchet distribution that is common to all households:

$$G(\varepsilon) = \exp\left\{\left(-\sum_{b \in B} \phi_b \varepsilon^{-\rho}\right)^{\frac{1}{\rho}}\right\} \tag{4.3}$$

where ρ is an inverse measure of the dispersion of households preferences over banks and ϕ_b captures amenities that make a bank more attractive to all consumers. Therefore, the average share of deposits across households held at bank b is given by:

$$\gamma_b = \frac{\phi_b R_{bd}^{\rho}}{\sum_{b'} \phi_{b'} R_{b'd}^{\rho}} \tag{4.4}$$

where R_{bd} is bank-specific deposit rate and ρ is the elasticity of deposit supply to banks. As is evident from the above equation, the share of deposits households place with a given bank depends positively on the bank's amenities (ϕ_b) and on its deposit rate (R_{bd}). The impact of deposit rates on household demand is determined by ρ: when $\rho < \infty$, banks face an upward sloping deposit supply but when $\rho \to \infty$, the deposit market is perfectly competitive such that banks offer the same deposit rate, up to a compensating differential.

The total household supply of deposits to a given bank is given by:

$$\mathscr{D}_b = \gamma_b w \tag{4.5}$$

Finally, the effective interest rate earned on deposits is given by:

$$R_{\mathscr{D}} = \sum_b \gamma_b R_{bd} \tag{4.6}$$

Firms

Production

There is a large number of non-financial firms $i \in N$, each of which produces a distinct product. Product markets are perfectly competitive.

Products can be either consumed by the household, or used as intermediate inputs for the production of other goods. The production function of firm i is a CES aggregate over value added and intermediate inputs:

$$y_i = (1-\mu)^{\frac{1}{\sigma}} (VA_i)^{\frac{\sigma-1}{\sigma}} + \mu^{\frac{1}{\sigma}} M_i^{\frac{\sigma-1}{\sigma}} \tag{4.7}$$

μ captures the share of intermediate inputs and σ captures the elasticity of substitution between value added and intermediate inputs.

Intermediate inputs are given by a CES aggregate over inputs sourced from all other firms:

$$M_i = \left[\sum_j a_{ji}^{\frac{1}{\zeta}} x_{ji}^{\frac{\zeta-1}{\zeta}} \right]^{\frac{\zeta}{\zeta-1}} \tag{4.8}$$

x_{ji} denotes the inputs sourced from firm j as intermediate inputs for the production of firm i. a_{ji} captures the importance of intermediates produced by firm j for the production of firm i. ζ is the elasticity of substitution between intermediates and final goods.

Value added is produced by mixing a continuum of intermediates, indexed by ω, which are aggregated via a CES function with elasticity of substitution η:

$$VA_i = \left[\int_0^1 [y_i(\omega)]^{\frac{\eta-1}{\eta}} d\omega \right]^{\frac{\eta}{\eta-1}} \tag{4.9}$$

Each intermediate good is produced with labour using linear technologies and a firm-wide productivity shifter:

$$y_i(\omega) = z_i l_i(\omega) \tag{4.10}$$

Financing

Firms trade intermediates at time 0, issuing perfectly enforceable, costless contracts that ensure payment at time 1 once final goods are sold. We assume a financing friction such that firms require financing in order to purchase labour. We follow Herreño (2023) and assume that for each variety ω, firms look for the cheapest financing option $R_i(\omega)$. Therefore, for each ω, firm i chooses from the following set:

$$R_i(\omega) = \min_{b \in B} \left\{ \frac{R_b}{\varepsilon_{ib}(\omega)} \right\} \tag{4.11}$$

where R_b is the effective interest offered by bank b to all firms. The cost to the firm is equal to the cost of funds R_b over an idiosyncratic shifter $\varepsilon_{ib}(\omega)$. The shifter is drawn randomly for each task-specific financing option and may reflect inherent differences in the nature of tasks or specific idiosyncrasies in the financing options offered by banks (e.g., a given bank may be better at financing task ω). The vector of draws $\varepsilon_{ib}(\omega)$ is drawn from a firm-specific Fréchet distribution:

$$F_i(\varepsilon) = \exp\left\{ \left(-\sum_{b \in B} T_{ib} \varepsilon_{ib}^{-\theta} \right)^{\frac{1}{\theta}} \right\} \tag{4.12}$$

where T_{ib} captures the absolute advantage of bank b at providing financing to firm i and θ captures the degree of substitutability of firms in switching banks. The Fréchet assumption, combined with the fact that all firms have to complete a unit interval of tasks, implies that the average share of credit financed by each bank b is given by:

$$s_{ib} = \frac{T_{ib} R_b^{-\theta}}{\sum_{b'} T_{ib'} R_{b'}^{-\theta}} \tag{4.13}$$

The above equation has intuitive implications. The share of credit that a firm obtains from a given bank depends positively on the firm-bank productivity shifter T_{ib}. It also depends negatively on the interest rate offered by the bank. The impact of interest rates on firm-level bank demand is determined by the elasticity of substitution θ: for higher θ, banks are more substitutable from the perspective of firms. Finally, the denominator captures the fact that the strength of the firm-bank relationship between a given firm i and a given bank b depends on the quality of that firm's relationships with all other banks, as well as the interest rates that these banks offer.

The credit demand by firm i is given by:

$$\mathscr{L}_{ib} = s_{ib} w l_i \tag{4.14}$$

Costs

Given the setup thus far, we can now consider the costs faced by firms. First, the total cost of producing a particular variety consists of the wage bill and the costs of financing that wage bill:

$$TC_i(\omega) = \frac{wl_i(\omega)}{z_i} R_i(\omega) \tag{4.15}$$

where R_i is the effective cost of bank credit for firm i and is given by

$$R_i = \left(\sum_b T_{ib} R_b^{-\theta} \right)^{-\frac{1}{\theta}} \tag{4.16}$$

Moreover, as a consequence of financing decisions, firm i's marginal cost of producing value added is

$$mc_i^{VA} = \frac{w}{z_i} R_i \tag{4.17}$$

which also denotes the unit price of labour. The marginal cost faced by firm i to produce output is given by

$$mc_i = \left[(1-\mu)(mc_i^{VA})^{1-\sigma} + \mu (P_i^M)^{1-\sigma} \right]^{\frac{1}{1-\sigma}} \tag{4.18}$$

where P_i^M is the intermediate input cost index of firm i.

Banks

Banks compete by choosing loan and deposit rates to maximise profit:

$$\pi_b = \mathscr{L}_b(R_b, .)R_b - \mathscr{D}_b(R_{bd}, .)R_{bd} \tag{4.19}$$

subject to a balance sheet constraint

$$\mathscr{L}_b = \sum_i \mathscr{L}_{ib} = \sum_i s_{ib} wl_i = \mathscr{D}_b + e_b \tag{4.20}$$

where e_b is the equity of the bank which is exogenous in our model, $\mathscr{L}_b(R_b,.)$ is the total lending by the bank, and $\mathscr{D}_b(R_{bd},.)$ is the total supply of deposits from households.

We assume the number of banks is large enough, so that banks treat the denominators capturing aggregate competition in the market for loans in (4.13) and in the market for deposits in (4.4) as parametric. Under these assumptions, banks have the same degree of market power in deposit and loan markets:

$$R_b = \frac{\frac{\theta}{\theta-1}}{\frac{\rho}{\rho+1}} R_{bd} \tag{4.21}$$

The optimal loan-deposit spread in (4.21) reflects imperfect competition in both deposit and loan markets. When deposit markets are perfectly competitive ($\rho \to \infty$), offered loan rates are simply equal to a markup over a common deposit rate (i.e., marginal cost). The level of the offered loan rates is determined by market clearing, which we discuss next.

Equilibrium

Firms price at marginal cost given by equation (4.18). Therefore,

$$p_i = \left[(1-\mu)(mc_i^{VA})^{1-\sigma} + \mu(P_i^M)^{1-\sigma}\right]^{\frac{1}{1-\sigma}} \tag{4.22}$$

where P_i^M is the intermediate input cost index of firm i and is given by

$$P_i^M = \left[\sum_j a_{ij} p_j^{1-\zeta}\right]^{\frac{1}{1-\zeta}} \tag{4.23}$$

The labour demand of firm i is given by

$$l_i = (1-\mu)\left(\frac{mc_i^{VA}}{p_i}\right)^{-\sigma} y_i \tag{4.24}$$

while the intermediate input demand of firm i is given by

$$x_{ji} = \mu a_{ji} \left(\frac{p_i}{P_i^M} \right)^{-\zeta} y_i \tag{4.25}$$

Bank profits are redistributed to households at time 1, so that consumption demand for firm i's product is given by:

$$c_i = \beta_i \frac{Y}{p_i} \tag{4.26}$$

where disposable income is $Y = R_{\mathscr{D}} w + \Pi$ and $\Pi = \sum_b \pi_b$ denotes aggregate bank profits.

Product market clearing requires that for each firm i, total production covers the final demand by consumer as well as the intermediate goods demand by other firms:

$$c_i + \sum_j x_{ji} = y_i \tag{4.27}$$

Financial market clearing requires that for each bank b, total loans equals total deposits plus equity

$$\sum_i s_{ib} w l_i = \mathscr{D}_b + e_b \tag{4.28}$$

We can now define an equilibrium.

Definition 1 *An equilibrium consists of prices for goods* $\mathbf{p}$*, labour* w*, loans* $\mathbf{R}$*, and deposits* $\mathbf{R_d}$ *such that:*

1. *Households maximise utility taking prices and deposit rates as given: bank deposits satisfy* (4.5) *and consumption demand satisfies* (4.26).

2. *Firms maximise profits taking prices and loan rates as given: prices satisfy* (4.22)*, financing decisions satisfy* (4.13)*, labour demands satisfy* (4.24)*, and input demands satisfy*

(4.25).

3. *Banks maximise profits: loan and deposit rates satisfy* (4.21).

4. *Product markets and financial markets clear:* (4.27) *holds for all firms and* (4.28) *holds for all banks.*

4.2.2 Impact of Bank-Specific Credit Supply Shocks

In this section, we analyse the model-implied interactions between financial and production networks, and how these shape the real effects of bank-specific lending shocks. This highlights how firms' model-implied responses to idiosyncratic shocks to bank's credit supply propagate to firms through the network of bank-firm relationships, described by s_{ib}, and how these shocks propagate through the production network. Overall, we are interested in the aggregate impact of these bank-specific shocks.

However, before examining the impact of bank-specific shocks to GDP, it is useful to first document how credit supply ("cs") shocks, in general, affect prices in our model. These could be shocks to either ϕ_b or e_b. We show that the impact on prices of such shocks depends on the both the bank-firm network and production network, which is informative about the impact of cs shocks on overall GDP.

To aid intuition, we consider an economy in which firms operate Cobb-Douglas technologies, that is $\sigma = 1$ and $\zeta = 1$. The equilibrium elasticity of the response of a firm's price p_i in response to a set of cs shocks , or more specifically, idiosyncratic shocks to bank b's funding, $d\log\phi_b$ or $d\log e_b$, is given by

$$\frac{d\log p_i}{d\log\phi_b} = \underbrace{(1-\mu)s_{ib}\frac{d\log R_b}{d\log\phi_b}}_{\substack{\text{direct impact of cs shock}\\\text{on firm borrowing cost}}} + \underbrace{(1-\mu)\sum_{b'\neq b} s_{ib'}\frac{d\log R'_b}{d\log\phi_b}}_{\substack{\text{indirect impact of cs shock}\\\text{on firm borrowing cost}}} + \underbrace{\mu\sum_j a_{ji}\frac{d\log p_j}{d\log\phi_b}}_{\substack{\text{downstream propagation}\\\text{of cs shock through eqm}\\\text{changes in supplier prices}}} \quad (4.29)$$

The first term on the right-hand side of (4.29) captures the direct impact of credit supply shocks on a firm's borrowing costs. This is increasing in s_{ib} which shows that the more reliant firm i is on bank b, the larger the direct impact. The second term captures the indirect impact of credit supply shocks on a firm's borrowing costs. This captures the equilibrium adjustment in financial markets through changes in the effective interest rates offered by all other banks. Finally, the third term captures the downstream propagation of the credit supply shock through equilibrium changes in prices of firm i's direct suppliers.

Next, we can stack the equilibrium responses across all firms in order to show that changes in prices depend on changes in equilibrium interest rates, weighted by the strength of the firm-bank relationships as well as the production network. Therefore, consider the following notation: $\hat{p}_b$ represents the vector of all equilibrium price responses of all firms to a credit supply shock, $\hat{R}_b$ represents the stacked vector of equilibrium responses of all banks' interest rates to the credit supply shock, $S = \{s_{ib}\}_{ib}$ represents the matrix of firm-bank linkages, and $A = \{a_{ij}\}_{ij}$ represents the input-output network. Given this notation, prices respond to changes in loan rates according to the following:

$$\hat{p}_b = (1-\mu)(1-\mu A')^{-1} S \hat{R}_b \tag{4.30}$$

Equation (4.30) captures a key insight of this paper: the bank-firm network S and the production network A jointly determine how shocks to banks' credit supply propagate to the real economy through product prices. Interest rates are themselves equilibrium objects. As a consequence, (4.30) also shows that the response of real outcomes to shocks to prices originating in the real economy (e.g., idiosyncratic productivity shocks to firms) depends on both production and bank-firm networks. Intuitively, as real shocks reallocate production across firms, this may affect interest rates through a reallocation of the demand for loans across banks. The subsequent changes in interest rates feed back into changes in prices, and so on.

With the understanding of how cs shocks affect prices, we now characterise the general equilibrium real effects of a set of shocks to banks' equity $\hat{\mathbf{e}} = [d\log e_b]_b$, assuming a Cobb-Douglas production structure.

Proposition 8 ***The aggregate impact of granular bank lending shocks***

Suppose that $\sigma = \zeta = 1$. Then the change in real GDP in response to idiosyncratic shocks to banks' equity $\hat{\mathbf{e}} = [d\log e_1, ..., d\log e_B]$, is given by:

$$d\log Y = -\left[(1-\omega_{\Pi}^{Y})(\mu\lambda' S - \gamma') - \omega_{\Pi}^{Y}\iota'\right]\mathscr{M}\left[\Gamma\alpha_b^e - \omega_{\Pi}^{Y} I_{B\times 1}\iota'\right]\hat{\mathbf{e}} + \omega_{\Pi}^{Y}\iota'\hat{\mathbf{e}}^1 \tag{4.31}$$

where $\omega_{\Pi}^{Y} = 1 - \frac{\frac{\theta}{\theta-1}}{\frac{\rho}{\rho+1}}\frac{wR_{\mathscr{D}}}{Y}$ denotes the share of bank profits in aggregate GDP, $\lambda = [\frac{p_i y_i}{Y}]_f$ is the vector of Domar weights, $S = [s_{ib}]_{i,b}$ is the network of bank-firm linkages, $\gamma = [\gamma_b]$ is the vector of banks' market shares in the deposit market, $I_{\alpha^e} \equiv Diag(\frac{e_b}{\mathscr{D}_b + e_b})$, and $\iota = \{\frac{e_b R_b}{\sum e_b R_b}\}_b$ is the vector of banks' normalised equity returns. Γ is a $B\times B$ matrix that captures interest rate spillovers across banks in partial equilibrium, $\Gamma_{b,b'} = \frac{\partial \ln R_b}{\partial \ln R_{b'}}$,

$$\Gamma = \left[-\theta I + (\theta-1)\Omega S - \rho \times I_{\alpha_b^{\mathscr{D}}} + \rho\left(1_{B\times 1}\otimes\gamma'\odot(\alpha^e)'\right)\right]^{-1}, \tag{4.32}$$

where $\Omega = \{\frac{s_{ib}}{\sum_i s_{i,b}}\}_{i,b}$ is the network of firm-bank linkages scaled by the lending of bank b and $\alpha_b^e = 1 - \alpha_b^{\mathscr{D}}$ is bank b's leverage ratio. Finally, $\mathscr{M} \equiv \left\{1 + (1-\alpha^e)\frac{\partial \ln w}{\partial \ln R_{b'}} - \frac{\partial \ln Y}{\partial \ln R_{b'}}\right\}_{b,b'\in\mathscr{B}}$ is a matrix of multipliers that capture the equilibrium feedback between interest rates and banks' loan demand and deposit supply:

$$\mathscr{M} \equiv \left[\Gamma^{-1} + \left(I_{1-\alpha^e} + \left(1-\omega_{\Pi}^{Y}\right)\right)\mu\lambda' S + I_{B\times 1}\left((1-\omega_{\Pi}^{Y})\gamma' + \omega_{\Pi}^{Y}\iota'\right)\right] \tag{4.33}$$

Proof: See Appendix D.1

The proposition makes clear the interplay of production and financial networks for the

[1] While this result is specific to a model with one factor of production (labour), it can be expanded to allow for multiple factors. While the model intuition will remain the same, the key difference with additional factors is that the impact of equity shock will increase in the number of factors that require external financing.

aggregate effects of granular bank credit supply shocks. Specifically, the general equilibrium change in GDP in response to equity shocks comprises two main effects: The first effect is given by the last term in (4.31) and operates through equity returns. Specifically, it captures that an increase in equity directly raises GDP by raising banks' equity returns, holding interest rates and factor income fixed.

The second effect is given by the first term in (4.31) and operates through adjustment in interest rates. Adjustments in interest rates directly impact GDP through financial profits, and indirectly through changing real output demand and hence wages. The second effect is subject to a multiplier, capturing how real wages and profits feed back into interest rates by shifting banks' deposit supply and credit demand.

An additional feature of our proposition is that banking sector competition (governed by θ and ρ) plays a role in both the first and second effects mentioned above. This is because it affects financial returns as well as the pass-through matrix Γ. This feature of our model speaks to an interesting literature on the role of banking sector competition in financial stability (e.g., Corbae and Levine (2024)) and Benchimol and Bozou (2024)). These papers typically debate the following trade-off: competition boosts market efficiency but by reducing banking sector profits, it worsens financial stability. We add another layer to this discussion: competition in the banking sector also affects the propagation of shocks.

One key feature of Proposition 8 is that we show that the aggregate effects of bank lending shocks are characterised by sufficient statistics that can be measured or estimated in the data. These sufficient statistics are the structure of the bank-firm network, the Domar weights, the share of intermediate inputs in production, the share of bank profits in aggregate GDP, bank equity returns, bank substitutability in loan and deposit markets, market shares of banks, and a bank's leverage ratio. Having such measurable sufficient statistics is important as it allows our theoretical framework to inform future empirical exercises. In the next section, we describe an example of an empirical exercise that one could use to complement our theoretical framework.

4.3 Possible Empirical Application

In this section, we describe a possible empirical application in the UK which, when implemented, would complement our theoretical framework.[2] In as much detail as is feasible without the data, we explain which data is necessary and their sources, possible measures of an exogenous credit supply shock, and the empirical strategy.

4.3.1 Data

Our theoretical framework highlights the data require to estimate an empirical complement to the model. In this section, we describe the data sources.

The bank-firm network (S) can be constructed using the Bureau van Dijk FAME database. This database provides annual accounts data on around 1.2 million UK registered companies. Importantly, it includes data on a bank-firm relationship if collateral was exchanged as part of the loan. Therefore, one can construct for each year t and any firm-bank pair (i,b), the stock of and change in the number of outstanding loans $l_{ib,t}$ held by firm i. Moreover, the annual accounts data also provides information on firm performance. Therefore, one can use this to obtain data on employment, sales, as well as gross value added.

Next, to empirically account for firm-firm network (A), one can utilise sector-level input-output tables provided by the UK Office of National Statistics (ONS). This can be used to trace input-output linkages between 104 industries and construct the Domar weights (λ). One can also use data from ONS to construct the shares of intermediate inputs used (μ).

The remaining sufficient statistics are related to banks. Bank leverage ratios (α^e) can be obtained from bank balance sheets as can be the market shares in deposits (γ) and bank equity returns (ι). The share of bank profits in aggregate GDP can be obtained from the ONS (ω_Π). Finally, estimates of substitutability in loan markets (θ) and deposit markets (ρ) can be found in Herreño (2023).

[2] In order to specifically estimate the model with data, one would need to extend our theoretical model to include dynamics.

Armed with measures of the sufficient statistics necessary for estimating the impact of granular bank lending shocks on GDP, one remaining challenge, covered in the next section, is identifying such exogenous shocks.

4.3.2 Granular Credit Supply Shocks

We suggest two different shocks that can be used as measures of exogenous credit supply shocks. While we focus primarily on the shock identified in Tracey and Sowerbutts (2023), we will also briefly describe the approach of Amiti and Weinstein (2018).[3]

Tracey and Sowerbutts (2023) introduce the idea of using misconduct provisions as an instrument for bank-specific credit supply shocks. The UK setting is intriguing for at least two reasons: Misconduct provisions were sizeable, and the differences in timing at which they occurred relative to when the misconduct took place are helpful for identification.

Between 2012 to 2015, UK banks saw a dramatic rise in misconduct provisions, primarily driven by the Payment Protection Insurance (PPI) misselling scandal. PPI covers loan repayments in specific circumstances and, in many cases, were sold to customers without their knowledge. Banks had strong incentives to sell PPI as they were very profitable products (often more profitable than the loan itself). While issues pertaining to PPI were known from the late 1990s, they continued to be sold. It was not until 2009 that PPI was banned, and an unexpected High Court ruling allowed new conduct standards on PPI to be applied retroactively. As a result of the High Court rule, UK banks provisioned around £87 billion, a sum that accounts for a meaningful share of aggregate capital and reflects the largest financial services redress exercise in UK history.

Data on these misconduct provisions can be found in the notes to the financial statements of banks' published accounts. From these accounts, one can also collect data on total lending, total deposits, total assets, as well as total equity.

An alternative, and commonly used, approach to construct bank-level credit supply

[3] One concern with using the shock in Tracey and Sowerbutts (2023) is that their results currently suggest negative shocks to equity may increase lending.

shocks is the method in Amiti and Weinstein (2018). In this method, a credit supply factor is identified for each bank as the bank fixed effect at a bank-firm-level regression of credit growth on bank- and firm-fixed effects. The firm-specific credit shocks are then calculated as the weighted average of the bank-specific factors, where the weights are a firm's pre-period exposure to each bank. The pre-period exposure is typically prior to a large shock (e.g., the global financial crisis).

Given the Amiti and Weinstein (2018) approach is widely used in the literature (e.g., Alfaro et al. (2021) and Huremović et al. (2023)), we focus on how one might use the first approach for constructing bank-level and firm-level credit supply shocks. Once the shocks are defined, one can use either approach in the empirical strategy we describe.

Bank-Level Credit Supply Shocks

Using the approach in Uppal (2024), one can translate unexpected changes in provisions to unexpected changes in bank equity making use of accounting identities. These changes in bank equity are what is needed to map our theoretical model to the data.

The first step is to document the impact of unexpected PPI provisions on bank profits for each bank using the following accounting identity:

$$\frac{\text{Profits}_{bt}}{\text{Assets}_{bt}} = \frac{\text{Net Interest Income}_{bt}}{\text{Assets}_{bt}} + \frac{\text{Net Noninterest Income}_{bt}}{\text{Assets}_{bt}} + \frac{\text{Other Income}_{bt}}{\text{Assets}_{bt}} - \frac{\text{PPI Provisions}_{bt}}{\text{Assets}_{bt}} - \frac{\text{Other Provisions}_{bt}}{\text{Assets}_{bt}} \tag{4.34}$$

The above will provide a method to explain the amount of variation in profits that is driven by the unexpected change in PPI provisions for a bank b at time t. The second step is to document how the unexpected change in profits (due to PPI provisions) leads to changes in equity. One can use the following accounting identity:

$$\frac{\text{Cumulative Profits}_{t}}{\text{Assets}_{t}} - \frac{\text{Cumulative Dividends}_{t}}{\text{Assets}_{t}} \approx \frac{\text{Equity}_{t}}{\text{Assets}_{t}} \tag{4.35}$$

This will provide a measure of unexpected changes in bank equity that are driven by PPI provisions. We call these bank-level credit supply shocks S_{bt}.

The key identifying assumption required for S_{bt} to constitute a valid credit supply shock is that misconduct provisions are uncorrelated with unobserved time-variant factors that influence lending behaviour of firms. As argued by Tracey and Sowerbutts (2023), the difference between "misconduct taking place" and "misconduct provision being made" is key to the underlying identification argument. All PPI provisions incurred by UK banks after 2010 were made in relation to fraudulent sales of PPI before 2009, with two thirds of PPI sales being made in the early 2000s or earlier, according to the UK Commission on Banking Standards.

Firm-Level Credit Supply Shocks

One can leverage heterogeneity in the strength of linkages in the network generated by firm-bank linkages through credit flows to construct a shift-share instrument for firm-level credit supply. Denote by n_{bit} the number of outstanding loans held by firm i in year t issued by bank b. n_{bit} provides a natural benchmark to measure the strength of firm-bank connections along the extensive margin.[4]

The shock to the credit supply of firm i in year t is denoted S_{it} and defined as a firm-specific average over the bank-level credit supply shocks experienced by all banks $b \in \mathscr{B}$:

$$S_{it} = \sum_{b \in \mathscr{B}} \omega_{ib,t_0} S_{bt} \tag{4.36}$$

where the firm-specific weight ω_{ib,t_0} is a shift share variable, corresponding to the share of firm i's total loans that bank b accounted for in a base year t_0:

$$\omega_{ib,t_0} = \frac{n_{bit}}{\sum_{b \in \mathscr{B}} n_{bit}} \tag{4.37}$$

[4]Note that we suggest the extensive margin as the FAME database only provides information on the number of loans.

The base year t_0 must lie outside of the baseline sample period. This approach is common in the literature using shift-share instruments (Bartik (1991)). In our context, it reflects the idea that past firm-bank relationships indicate firm-specific, unobserved factors that cause a firms' overall supply of credit to depend more on a specific bank. For example, if a firm exclusively deals with one specific bank, this could indicate that that bank has an unobserved advantage in providing funding to said firm. As a result, shocks to the lending capacity of that bank should have a disproportionate effect on the total amount of funding available to that firm. Goldsmith-Pinkham et al. (2020) argue that as long as past shift shares are not correlated with current unobserved changes also affecting the total credit available to i, then (4.36) provides a valid instrument for firm-level credit supply. It is worth noting that, thus, for (4.36) to constitute a firm-level credit supply shock, banks have to be imperfect substitutes in terms of their ability to provide funding from the perspective of individual firms.

Structural Interpretation

Recent work by Herreño (2023) provides a framework to analyse the formation of bank-firm relationships, which help illuminate much of the above discussion and how it pertains to the logic of the shift-share instrument. Our model features similar forces, and therefore the implied firm-level effects from bank-specific credit supply shocks are conceptually equivalent to (4.36).

Banks compete to provide funding to firms, and from the perspective of firms, banks are heterogeneous in terms of their capacity to provide funding. Using a probabilistic formulation following Eaton and Kortum (2002) to characterise the funding choice problem, the share of financing provided by a bank b is given by:

$$s_{ib} = \frac{T_{ib} R_b^{-\theta}}{\sum_{b \in \mathscr{B}} T_{ib'} R_{b'}^{-\theta}} \tag{4.38}$$

This is precisely the share we have in our theoretical model (equation (4.13)) where T_{ib} captures the exogenous strength of the banking relationship between firm i and bank b, and R_b denotes the effective interest rate offered by bank b. Note that R_b is not firm-specific, and thus it can be thought of as encapsulating bank-specific credit supply shocks. The overall cost of external credit, which is also an inverse measure of the effective credit supply available to firm i, is denoted R_i and depends on quality of its overall firm-bank relationships, the effective rates offered by banks, as well as the elasticity of substitution between banks which is captured by θ:

$$R_i = \left(\sum_{b \in \mathscr{B}} T_{ib} R_b^{-\theta} \right)^{-\frac{1}{\theta}} \tag{4.39}$$

In this setting, the change in a firm's overall cost of funding, $d\log R_i$, in response to a set of bank-specific credit supply shocks, $d\log R_b$, that are unrelated with potential changes in T_{ib} can be expressed as:

$$d\log R_i = \sum_{b \in \mathscr{B}} s_{ib} d\log R_b \tag{4.40}$$

Equation (4.40) shows that bank-specific credit supply shocks in the model can be "translated" into firm-specific credit supply shocks. The appropriate weights, given by (4.38), are sufficient statistics summarising the importance of bank b as a credit supplier for firm i. The shares s_{ib} in (4.38) capture unobserved heterogeneity in firm-bank specific relationship efficiencies T_{ib} as well as the imperfect substitutability of banks. Lastly, it is worth noting that (4.40) only appropriately describes the firm-specific impact of bank-specific shocks if these shocks are uncorrelated with any unobserved changes in T_{ib}.

In summary, appealing to a structural microfoundation highlights the following assumptions are sufficient for identification: First, firms cannot costlessly switch between banks to serve their funding needs, so banks are imperfect substitutes. Second, the weights in (4.37) summarise time-invariant factors determining the observed heterogeneity in firm-bank linkages, and are not correlated with future changes in these factors. Third, the bank-specific shocks are not correlated with unobserved shocks to firm-bank specific credit demand (T_{ib}).

We also note that one can reinterpret shift-share instruments as shock-level regressions. Borusyak et al. (2022) provide alternative conditions for consistency based on this reinterpretation. As we believe that the identification assumptions yield a priori plausibly exogenous shocks at the bank-level, one can also implement the tools provided by the Borusyak et al. (2022) to assess the ex-post plausibility of the assumption of exogenous shocks.

4.3.3 Empirical Approach

In this section, we describe an approach to estimating (i) the direct effect of firm-level credit supply shocks resulting from policy driven shocks to bank lending, and (ii) the degree to which these shocks propagate through the input-output network. For the latter, we propose two empirical approaches, following Alfaro et al. (2021) and Carvalho et al. (2021) respectively.

4.3.4 Direct Effects

To estimate the direct effect of credit supply shocks on firm outcomes, one can estimate the following regressions:

$$y_{it} = \alpha S_{it} + \beta \mathbf{X}_{it} + FE + \varepsilon_{it} \tag{4.41}$$

where y_{it} denotes either the log of firm sales or employment. $\mathbf{X}_{it}$ is a vector of time-varying firm characteristics, including economic and financial variables. FE denotes a set of fixed effects comprising either a set of sector/year, sector $\times$ year or firm and year dummies.

The coefficient of interest is α. It captures how differences in firm-level exposure to bank-level credit supply shocks affect the relative economic performance of firms, as measured by either employment or sales growth. Relative to recent work by Alfaro et al. (2021) and Huremović et al. (2023) that relies on standard operationalisation of credit supply shocks following Amiti and Weinstein (2018) to estimate the real effects bank shocks, the estimated effect proposed here provides direct and transparent evidence on the transmission of financial shocks to the real economy through the firm-bank network by utilising unexpected changes in bank equity through PPI provisions.

4.3.5 Propagation

Understanding how idiosyncratic shocks to firms propagate through the input-output network is the subject of a large literature. This literature highlights that shocks affecting output and prices of one supplier may affect another firm by the way in which these two firms are connected through supplier and buyer relationships. While the majority of papers focusing on propagation is concerned with firm-specific productivity or demand shocks, Alfaro et al. (2021) and Huremović et al. (2023) provide important evidence on how credit supply shocks to Spanish firms propagate through the production network. However, neither paper explicitly models the banking sector which, as we show in Proposition 8, is critical to understand how shocks are propagated and amplified. Indeed, our model has specific sufficient statistics relating to measures of competition in the banking sector which would support estimating a fully quantified dynamic version of our theoretical model.

Here, we outline two approaches to uncovering production network propagation in the context of our setting. The first one closely follows Di Giovanni et al. (2018) and has the advantage that it can be estimated from cross-sectional regressions similar to (4.41), but it accounts for indices proxying for a firms' exposure to shocks originating in the rest of the economy. The second builds on the empirical approach in Carvalho et al. (2021). While this approach, in principle, provides more direct evidence on propagation, it is substantially more data-intensive.

Therefore, we suggest building on Di Giovanni et al. (2018) by using the sectoral Input-Output (IO) tables for the UK in conjunction with firms' sectoral identifiers provided in the FAME database to construct indices that aim to capture the effects of downstream and upstream propagation. One can combine information on sectoral linkages from IO tables with information on firms' wage bills relative to sales and construct proxies for a firm's intermediate input intensity:

$\omega^I_{it} = 1 - \frac{\text{Wage Bill}_{it}}{\text{Sales}_{it}}$. Similarly, ω^S_{it} denotes the market share of firm i active in sector j relative to total sales in sector j. The matrix A documents sectoral direct input-output linkages, where a_{jk} denotes the cost-share of intermediates sourced from sector j in sector k's total usage. The Leontief matrix $L = (I - A)^{-1}$ captures direct and all indirect linkages between two sectors. Specifically entry l_{jk} captures the importance of sector j as a supplier to sector k, taking into account all direct and indirect paths that connect these two sectors. Finally, denote by Σ_{kt} the sales weighted average of the firm-specific credit supply shocks S_{it} of all firms belonging to sector k.

Using these data points, we can define firm-specific measures of downstream propagation. Down^1_{it} measures the effect of shocks received via direct linkages, while Down^∞_{it} captures how shocks to suppliers of a firm i active in a given sector j proxy for the impact of propagation through all possible paths.

$$\text{Down}^1_{it,j} = \omega^I_{it,j} \sum_k a_{kj} \Sigma_{kt} \qquad \text{Down}^\infty_{it,j} = \omega^I_{it,j} \sum_k l_{kj} \Sigma_{kt} \tag{4.42}$$

Firm-specific measures of upstream propagation, in turn, aim to proxy for shocks received by firm i in sector j from its customers.

$$\text{Up}^1_{it,j} = \omega^S_{it,j} \sum_k a_{jk} \Sigma_{kt} \qquad \text{Up}^\infty_{it,j} = \omega^S_{it,j} \sum_k l_{jk} \Sigma_{kt} \tag{4.43}$$

These measures of propagation allow one to estimate the relative importance of direct effects versus indirect effects from either upstream or downstream propagation:

$$y_{it} = \alpha S_{it} + \alpha_1^{(1)} \text{Down}^1_{it,j} + \alpha_2^{(1)} \text{Up}^1_{it,j} + \alpha_1^{(\infty)} \text{Down}^\infty_{it,j} + \alpha_2^{(\infty)} \text{Up}^\infty_{it,j} + \beta \mathbf{X}_{it} + FE + \varepsilon_{it} \tag{4.44}$$

The results of estimating the empirical model in (4.44) inform the relative magnitude of direct and indirect channels through which credit supply shocks affect firms. These empirical results would complement a small but growing literature estimating the propagation of financial shocks through production networks.

4.4 Conclusion

In this paper, we show the importance of accounting for the structure of both financial and production networks in order to understand the channels through which financial shocks transmit to the real economy. We do so by developing a general equilibrium model featuring production networks, firm-bank relationships, and financial frictions. We tractably introduce a financial sector with heterogeneous banks into a benchmark model of production networks. The model highlights new mechanisms through which idiosyncratic shocks to bank lending permeate through financial and production networks.

We show that the aggregate effects on output from a bank-level credit supply shock depend not only on the production network, but also on how easily firms can switch between banks. Specifically, we document that the general equilibrium propagation of micro-level shocks depends on structural parameters governing competition in the banking sector as well as the leverage of banks. This additional propagation essentially amplifies the direct effects of credit supply shocks which highlights the importance of accounting for network interconnections.

Moreover, we show that the aggregate effects of bank-level shocks can be characterised by a set of sufficient statistics that can be measured or estimated in the data. These include the structure of the bank-firm network, the Domar weights, the share of intermediate inputs in production, the share of bank profits in GDP, bank equity returns, bank substitutability in loan and deposit markets, market shares of banks, and a bank's leverage ratio. This characterisation lends our framework to empirical study.

Therefore, we suggest an empirical setting exploits unexpected misconduct fines for UK banks in line with Tracey and Sowerbutts (2023). We show how these bank-level credit supply shocks can be translated in firm-level credit supply shocks and therefore used to estimate the importance of direct versus indirect effects, following the approach of Di Giovanni et al. (2018).

Specifically, we show that these instruments have a structural interpretation in line with our theoretical model. We believe our theoretical framework provides an avenue for both future empirical work, as highlighted above, but also a basis for developing a dynamic model that can be used for quantification.

4.5 Chapter Acknowledgements

Chapter 4 is currently being planned for submission for publication of the material. Fabian Trottner and the dissertation author, Ali Uppal, are co-authors of this chapter.

Bibliography

Acemoglu, D., Carvalho, V. M., Ozdaglar, A., and Tahbaz-Salehi, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016.

Adrian, T., Boyarchenko, N., and Shin, H. (2019). On the Scale of Financial Intermediaries. *Federal Reserve Bank of New York Staff Reports*, 56(10):25–27.

Adrian, T., Moench, E., and Shin, H. S. (2014). Dynamic leverage asset pricing. *Federal Reserve Bank of New York Staff Reports*, (625).

Adrian, T. and Shin, H. S. (2010). Liquidity and leverage. *Journal of financial intermediation*, 19(3):418–437.

Afonso, G., Santos, J. A., and Traina, J. (2014). Do "too-big-to-fail" banks take on more risk? *Economic Policy Review*, 20(2):41–58.

Aharony, J., Saunders, A., and Swary, I. (1986). The effects of a shift in monetary policy regime on the profitability and risk of commercial banks. *Journal of Monetary Economics*, 17(3):363–377.

Ai, H. and Bansal, R. (2018). Risk preferences and the macroeconomic announcement premium. *Econometrica*, 86(4):1383–1430.

Ajello, A., Boyarchenko, N., Gourio, F., and Tambalotti, A. (2022). Financial stability considerations for monetary policy: Theoretical mechanisms. *FRB of New York Staff Report*, (1002).

Alessandri, P. and Nelson, B. D. (2015). Simple banking: Profitability and the yield curve. *Journal of Money, Credit and Banking*, 47(1):143–175.

Alfaro, L., García-Santana, M., and Moral-Benito, E. (2021). On the direct and indirect real effects of credit supply shocks. *Journal of Financial Economics*, 139(3):895–921.

Allen, F. and Gale, D. (2000). Financial contagion. *Journal of political economy*, 108(1):1–33.

Altavilla, C., Boucinha, M., and Peydró, J.-L. (2018). Monetary policy and bank profitability in a low interest rate environment. *Economic Policy*, 33(96):531–586.

Altavilla, C., Brugnolini, L., Gürkaynak, R. S., Motto, R., and Ragusa, G. (2019). Measuring euro area monetary policy. *Journal of Monetary Economics*, 108:162–179.

Amiti, M. and Weinstein, D. E. (2018). How much do idiosyncratic bank shocks affect investment? evidence from matched bank-firm loan data. *Journal of Political Economy*, 126(2):525–587.

Ampudia, M. and van den Heuvel, S. J. (2019). Monetary Policy and Bank Equity Values in a Time of Low and Negative Interest Rates. Finance and Economics Discussion Series 2019-064, Board of Governors of the Federal Reserve System (U.S.).

Angeloni, I. and Faia, E. (2013). Capital regulation and monetary policy with fragile banks. *Journal of Monetary Economics*, 60(3):311–324.

Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review*, 109(6):2333–2367.

Balloch, C. and Koby, Y. (2023). Low rates and bank loan supply: Theory and evidence from japan. *Working Paper*.

Baqaee, D. R. (2018). Cascading failures in production networks. *Econometrica*, 86(5):1819–1838.

Baqaee, D. R. and Farhi, E. (2019). The macroeconomic impact of microeconomic shocks: Beyond hulten's theorem. *Econometrica*, 87(4):1155–1203.

Barro, R. J. and Gordon, D. B. (1983). Rules, discretion and reputation in a model of monetary policy. *Journal of Monetary Economics*, 12(1):101–121.

Barrot, J.-N. and Sauvagnat, J. (2016). Input specificity and the propagation of idiosyncratic shocks in production networks. *The Quarterly Journal of Economics*, 131(3):1543–1592.

Bartik, T. J. (1991). *Who benefits from state and local economic development policies?* W.E. Upjohn Institute for Employment Research.

Begenau, J., Bigio, S., Majerovitz, J., and Vieyra, M. (2021). A q-theory of banks. *National Bureau of Economic Research Working Paper Series*, No. 27935.

Begenau, J. and Landvoigt, T. (2021). Financial regulation in a quantitative model of the modern banking system. *National Bureau of Economic Research Working Paper Series*, No. 28501.

Benchimol, J. and Bozou, C. (2024). Desirable Banking Competition and Stability. *Working Paper*.

Bernanke, B. S. (2015). Should monetary policy take into account risks to financial stability? *Brookings Commentary*, https://www.brookings.edu/articles/should-monetary-policy-take-into-account-risks-to-financial-stability/.

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

Bernanke, B. S., Gertler, M., and Gilchrist, S. (1999). The financial accelerator in a quantitative business cycle framework. *Handbook of Macroeconomics*, 1:1341–1393.

Bianchi, F., Gómez-Cram, R., Kind, T., and Kung, H. (2023). Threats to central bank independence: High-frequency identification with twitter. *Journal of Monetary Economics*, 135:37–54.

Bigio, S. and La'o, J. (2020). Distortions in production networks. *The Quarterly Journal of Economics*, 135(4):2187–2253.

Blinder, A. S., Ehrmann, M., Fratzscher, M., De Haan, J., and Jansen, D.-J. (2008). Central bank communication and monetary policy: A survey of theory and evidence. *Journal of Economic Literature*, 46(4):910–45.

Boehm, C. E., Flaaen, A., and Pandalai-Nayar, N. (2019). Input linkages and the transmission of shocks: Firm-level evidence from the 2011 tōhoku earthquake. *Review of Economics and Statistics*, 101(1):60–75.

Borusyak, K., Hull, P., and Jaravel, X. (2022). Quasi-experimental shift-share research designs. *The Review of Economic Studies*, 89(1):181–213.

Boyarchenko, N., Favara, G., and Schularick, M. (2022). Financial Stability Considerations for Monetary Policy: Empirical Evidence and Challenges. *Federal Reserve Bank of New York Staff Reports*, No. 1003.

Bradley, D., Finer, D. A., Gustafson, M., and Williams, J. (2020). When Bankers Go to Hail: Insights into Fed-Bank Interactions from Taxi Data. *Working Paper*.

Bradley, D., Finer, D. A., Gustafson, M., and Williams, J. (2023). When bankers go to hail: Insights into fed–bank interactions from taxi data. *Management Science*.

Brusa, F., Savor, P., and Wilson, M. (2020). One Central Bank to Rule Them All. *Review of Finance*, 24(2):263–304.

Bu, C., Rogers, J., and Wu, W. (2021). A unified measure of fed monetary policy shocks. *Journal of Monetary Economics*, 118:331–349.

Busch, R. and Memmel, C. (2015). Banks net interest margin and the level of interest rates. Vfs annual conference 2015 (muenster): Economic development - theory and policy, Verein für Socialpolitik / German Economic Association.

Caballero, R. and Simsek, A. (2022). A Monetary Policy Asset Pricing Model. *National Bureau of Economic Research Working Paper Series*, No. 30132.

Caballero, R. J. and Simsek, A. (2020a). Prudential monetary policy. *National Bureau of Economic Research Working Paper Series*, No. 25977.

Caballero, R. J. and Simsek, A. (2020b). A risk-centric model of demand recessions and speculation. *The Quarterly Journal of Economics*, 135(3):1493–1566.

Calvo, D., Crisanto, J. C., Hohl, S., and Gutiérrez, O. P. (2018). Financial supervisory architecture: what has changed after the crisis? *FSI Insights on policy implementation*, 8.

Calza, A., Monacelli, T., and Stracca, L. (2013). Housing finance and monetary policy. *Journal of the European Economic Association*, 11(suppl_1):101–122.

Campbell, J. Y. and Cocco, J. F. (2015). A model of mortgage default. *The Journal of Finance*, 70(4):1495–1554.

Carvalho, V. M., Nirei, M., Saito, Y. U., and Tahbaz-Salehi, A. (2021). Supply chain disruptions: Evidence from the great east japan earthquake. *The Quarterly Journal of Economics*, 136(2):1255–1321.

Chodorow-Reich, G. (2014). The employment effects of credit market disruptions: Firm-level evidence from the 2008–9 financial crisis. *The Quarterly Journal of Economics*, 129(1):1–59.

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):1–45.

Cieslak, A., Morse, A., and Vissing-Jørgensen, A. (2019). Stock returns over the FOMC cycle. *The Journal of Finance*, 74(5):2201–2248.

Cieslak, A. and Vissing-Jørgensen, A. (2020). The Economics of the Fed Put. *The Review of Financial Studies*, 34(9):4045–4089.

Coibion, O. (2012). Are the effects of monetary policy shocks big or small? *American Economic Journal: Macroeconomics*, 4(2):1–32.

Corbae, D. and Levine, R. (2022). Competition, Stability, and Efficiency in the Banking Industry. *Unpublished working paper*, 8.

Corbae, D. and Levine, R. (2024). Competition, Stability, and Efficiency in the Banking Industry. *Working Paper*.

Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50(6):1431–1451.

Curdia, V. and Woodford, M. (2010). Credit spreads and monetary policy. *Journal of Money, credit and Banking*, 42:3–35.

Dell'Ariccia, G., Laeven, L., and Marquez, R. (2014). Real interest rates, leverage, and bank risk-taking. *Journal of Economic Theory*, 149(1):65–99.

Di Giovanni, J., Levchenko, A. A., and Mejean, I. (2018). The micro origins of international business-cycle comovement. *American Economic Review*, 108(1):82–108.

Di Tella, S. and Kurlat, P. (2021). Why Are Banks Exposed to Monetary Policy? *American Economic Journal: Macroeconomics*, 13(4):295–340.

Drechsler, I., Savov, A., and Schnabl, P. (2017). The Deposits Channel of Monetary Policy. *Quarterly Journal of Economics*, 132(4):1819–1876.

Drechsler, I., Savov, A., and Schnabl, P. (2018a). Liquidity, Risk Premia, and the Financial Transmission of Monetary Policy. *Annual Review of Financial Economics*, 10:309–328.

Drechsler, I., Savov, A., and Schnabl, P. (2018b). A model of monetary policy and risk premia. *The Journal of Finance*, 73(1):317–373.

Drechsler, I., Savov, A., and Schnabl, P. (2021). Banking on deposits: Maturity transformation without interest rate risk. *The Journal of Finance*, 76(3):1091–1143.

Eaton, J. and Kortum, S. (2002). Technology, geography, and trade. *Econometrica*, 70(5):1741–1779.

English, W. B., Van den Heuvel, S. J., and Zakrajšek, E. (2018). Interest rate risk and bank equity valuations. *Journal of Monetary Economics*, 98:80–97.

Faust, J. (2016). Oh What a Tangled Web We Have: Monetary Policy Transperency in Divisive Times. *Unpublished working paper*.

Finer, D. A. (2018). What insights do taxi rides offer into federal reserve leakage? *Chicago Booth: George J. Stigler Center for the Study of the Economy & the State Working Paper*, (18).

Flannery, M. J. and James, C. M. (1984). The effect of interest rate changes on the common stock returns of financial institutions. *The Journal of Finance*, 39(4):1141–1153.

Fostel, A. and Geanakoplos, J. (2008). Leverage cycles and the anxious economy. *American Economic Review*, 98(4):1211–44.

Freixas, X., Parigi, B. M., and Rochet, J.-C. (2000). Systemic risk, interbank relations, and liquidity provision by the central bank. *Journal of money, credit and banking*, pages 611–638.

Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica*, 79(3):733–772.

Galati, G. and Moessner, R. (2018). What Do We Know About the Effects of Macroprudential Policy? *Economica*, 85:735–770.

Galí, J. (1999). Technology, employment, and the business cycle: do technology shocks explain aggregate fluctuations? *American economic review*, 89(1):249–271.

Galí, J. (2015). *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*. Princeton University Press.

Gaudêncio, J., Mazany, A., and Schwarz, C. (2019). The impact of lending standards on default rates of residential real-estate loans. *ECB Occasional Paper*, (220).

Geanakoplos, J. (2010). The leverage cycle. *NBER macroeconomics annual*, 24(1):1–66.

Gertler, M. and Karadi, P. (2011). A model of unconventional monetary policy. *Journal of monetary Economics*, 58(1):17–34.

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

Gertler, M. and Kiyotaki, N. (2010). Financial intermediation and credit policy in business cycle analysis. In *Handbook of monetary economics*, volume 3, pages 547–599. Elsevier.

Gertler, M. and Kiyotaki, N. (2015). Banking, liquidity, and bank runs in an infinite horizon economy. *American Economic Review*, 105(7):2011–2043.

Gertler, M., Kiyotaki, N., and Prestipino, A. (2020). A macroeconomic model with financial panics. *The Review of Economic Studies*, 87(1):240–288.

Ghote, A. V. D. (2021). Interactions between monetary and macroprudential policy. *American Economic Journal: Macroeconomics*, 13(1):1–34.

Goldberg, J. and López-Salido, D. (2023). “Sowing the Wind” Monetary Policy. *Working Paper*.

Goldsmith-Pinkham, P., Sorkin, I., and Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624.

Gorodnichenko, Y., Pham, T., and Talavera, O. (2023). The voice of monetary policy. *American Economic Review*, 113(2):548–584.

Gourio, F., Kashyap, A. K., and Sim, J. W. (2018). The Trade offs in Leaning Against the Wind. *IMF Economic Review*, 66.

Gregory, A., Tharyan, R., and Christidis, A. (2013). Constructing and testing alternative versions of the Fama–French and Carhart models in the UK. *Journal of Business Finance & Accounting*, 40(1-2):172–214.

Grimm, M., Jordà, Ò., Schularick, M., and Taylor, A. M. (2023). Loose monetary policy and financial instability. *National Bureau of Economic Research Working Paper Series*, No. 30958.

Gürkaynak, R. S., Sack, B. P., and Swanson, E. T. (2007). Market-based measures of monetary policy expectations. *Journal of Business & Economic Statistics*, 25(2):201–212.

Guttman, I., Kadan, O., and Kandel, E. (2010). Dividend stickiness and strategic pooling. *The Review of Financial Studies*, 23(12):4455–4495.

Hansen, S., McMahon, M., and Tong, M. (2019). The long-run information effect of central bank communication. *Journal of Monetary Economics*, 108:185–202.

He, Z., Kelly, B., and Manela, A. (2017). Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics*, 126(1):1–35.

Heipertz, J. (2021). The banks along the river: Banking origins of aggregate fluctuations. *Working Paper*.

Hellwig, M. (1994). Liquidity provision, banking, and the allocation of interest rate risk. *European Economic Review*, 38(7):1363–1389.

Herreño, J. (2023). Aggregating the Effect of Bank Credit Supply Shocks on Firms. *Working Paper*.

Holmstrom, B. and Tirole, J. (1997). Financial intermediation, loanable funds, and the real sector. *the Quarterly Journal of economics*, 112(3):663–691.

Hu, G. X., Pan, J., Wang, J., and Zhu, H. (2022). Premium for heightened uncertainty: Explaining pre-announcement market returns. *Journal of Financial Economics*, 145(3):909–936.

Huremović, K., Jiménez, G., Moral-Benito, E., Peydró, J.-L., and Vega-Redondo, F. (2023).

Production and financial networks in interplay: Crisis evidence from supplier-customer and credit registers. *Working Paper*.

Jiang, E. X., Matvos, G., Piskorski, T., and Seru, A. (2023). Monetary tightening and u.s. bank fragility in 2023: Mark-to-market losses and uninsured depositor runs? Working Paper 31048, National Bureau of Economic Research.

Jiménez, G., Mian, A., Peydró, J.-L., and Saurina, J. (2020). The real effects of the bank lending channel. *Journal of Monetary Economics*, 115:162–179.

Jiménez, G., Ongena, S., Peydró, J.-L., and Saurina, J. (2012). Credit supply and monetary policy: Identifying the bank balance-sheet channel with loan applications. *American Economic Review*, 102(5):2301–2326.

Jiménez, G., Ongena, S., Peydró, J.-L., and Saurina, J. (2017). Macroprudential policy, countercyclical bank capital buffers, and credit supply: Evidence from the spanish dynamic provisioning experiments. *Journal of Political Economy*, 125(6):2126–2177.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American economic review*, 95(1):161–182.

Känzig, D. R. (2021). The macroeconomic effects of oil supply news: Evidence from opec announcements. *American Economic Review*, 111(4):1092–1125.

Kashyap, A. K. and Stein, J. C. (1994). Monetary policy and bank lending. In *Monetary policy*, pages 221–261. The University of Chicago Press.

Khwaja, A. I. and Mian, A. (2008). Tracing the impact of bank liquidity shocks: Evidence from an emerging market. *American Economic Review*, 98(4):1413–1442.

Kirti, D. (2020). Why do bank-dependent firms bear interest-rate risk? *Journal of Financial Intermediation*, 41:100823.

Kurov, A., Wolfe, M. H., and Gilbert, T. (2021). The disappearing pre-fomc announcement drift. *Finance research letters*, 40:101781.

Laarits, T. (2020). Pre-Announcement Risk. *NYU Stern School of Business*.

Laux, C. and Rauter, T. (2017). Procyclicality of US bank leverage. *Journal of Accounting Research*, 55(2):237–273.

Li, Z. (2022). Financial intermediary leverage and monetary policy transmission. *European Economic Review*, 144:104080.

Lucca, D. O. and Moench, E. (2015). The Pre-FOMC announcement drift. *Journal of Finance*, 70(1):329–371.

Maggiori, M. (2017). Financial intermediation, international risk sharing, and reserve currencies. *American Economic Review*, 107(10):3038–3071.

Martin, A., Mendicino, C., and Ghote, A. V. D. (2021). On the interaction between monetary and macroprudential policies. *ECB Working Paper Series*, No. 2527(2527).

Martinez-Miera, D. and Repullo, R. (2017). Search for yield. *Econometrica*, 85(2):351–378.

Martinez-Miera, D. and Repullo, R. (2019). Monetary Policy, Macroprudential Policy, and Financial Stability. *Annual Review of Economics*, 11:809–832.

Masciandaro, D., Peia, O., and Romelli, D. (2024). Central bank communication and social media: From silence to twitter. *Journal of Economic Surveys*, 38(2):365–388.

McLean, R. D. and Pontiff, J. (2016). Does academic research destroy stock return predictability? *The Journal of Finance*, 71(1):5–32.

Metrick, A. (2024). The failure of silicon valley bank and the panic of 2023. *Journal of Economic Perspectives*, 38(1):133–152.

Miranda-Agrippino, S. and Rey, H. (2020). US monetary policy and the global financial cycle. *The Review of Economic Studies*, 87(6):2754–2776.

Montiel Olea, J. L. and Plagborg-Møller, M. (2021). Local projection inference is simpler and more robust than you think. *Econometrica*, 89(4):1789–1823.

Morse, A. and Vissing-Jørgensen, A. (2020). Information transmission from the federal reserve to the stock market: Evidence from governors' calendars. *Working Paper*.

Nuño, G. and Thomas, C. (2017). Bank leverage cycles. *American Economic Journal: Macroeconomics*, 9(2):32–72.

Ottonello, P. and Song, W. (2022). Financial Intermediaries and the Macroeconomy: Evidence from a High-Frequency Identification. *National Bureau of Economic Research Working Paper Series*, No. 29638.

Pflueger, C., Siriwardane, E., and Sunderam, A. (2020). Financial market risk perceptions and the macroeconomy. *The Quarterly Journal of Economics*, 135(3):1443–1491.

Ramey, V. (2016). Macroeconomic Shocks and Their Propagation. volume 2 of *Handbook of Macroeconomics*, chapter 2, pages 71–162. Elsevier.

Rannenberg, A. (2016). Bank leverage cycles and the external finance premium. *Journal of Money, Credit and Banking*, 48(8):1569–1612.

Rigobon, R. and Sack, B. (2004). The impact of monetary policy on asset prices. *Journal of monetary economics*, 51(8):1553–1575.

Rogoff, K. (1985). The optimal degree of commitment to an intermediate monetary target. *Quarterly Journal of Economics*, 100.

Romer, C. D. and Romer, D. H. (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review*, 94(4):1055–1084.

Samuelson, P. A. (1945). The effect of interest rate increases on the banking system. *The American economic review*, 35(1):16–27.

Shapiro, A. H. and Wilson, D. J. (2022). Taking the fed at its word: A new approach to estimating central bank objectives using text analysis. *The Review of Economic Studies*, Forthcoming.

Sims, E. and Wu, J. C. (2021). Evaluating central banks' tool kit: Past, present, and future. *Journal of Monetary Economics*, 118:135–160.

Smets, F. (2014). Financial Stability and Monetary Policy: How Closely Interlinked? *International Journal of Central Banking*, 10(2).

Smets, F. and Wouters, R. (2007). Shocks and frictions in us business cycles: A bayesian dsge approach. *American economic review*, 97(3):586–606.

Stein, J. C. (1989). Cheap talk and the fed: A theory of imprecise policy announcements. *The American Economic Review*, 79(1):32–42.

Stein, J. C. and Sunderam, A. (2018). The fed, the bond market, and gradualism in monetary policy. *The Journal of Finance*, 73(3):1015–1060.

Stock, J. H. and Watson, M. W. (2018). Identification and estimation of dynamic causal effects in macroeconomics using external instruments. *The Economic Journal*, 128(610):917–948.

Svensson, L. E. O. (2017). Cost-benefit analysis of leaning against the wind. *Journal of Monetary Economics*, 90:193–213.

Svensson, L. E. O. (2018). Monetary policy and macroprudential policy: Different and separate? *Canadian Journal of Economics*, 51(3):802–827.

Taylor, J. B. (1993). Discretion versus policy rules in practice. In *Carnegie-Rochester conference series on public policy*, volume 39, pages 195–214. Elsevier.

Tinbergen, J. (1952). *On the Theory of Economic Policy*. Amsterdam: North-Holland.

Tracey, B. and Sowerbutts, R. (2023). How Do Banks Respond to Capital Shocks? Evidence from Misconduct Provisions. *Working Paper*.

Uppal, A. (2024). Do Higher Interest Rates Make The Banking System Safer? Evidence From Bank Leverage. *Unpublished working paper*.

Van den Heuvel, S. (2009). The bank capital channel of monetary policy. *Working Paper*.

Vissing-Jorgensen, A. (2020). Central Banking with Many Voices: The Communications Arms Race. *Unpublished working paper*.

Walter, J. R. (1991). Loan loss reserves. *FRB Richmond Economic Review*, July/August:20–30.

Wieland, J. F. and Yang, M.-J. (2020). Financial Dampening. *Journal of Money, Credit and Banking*, 52(1):79–113.

Wolf, C. K. and McKay, A. (2023). What Can Time-Series Regressions Tell Us About Policy Counterfactuals? *Econometrica*, Forthcoming.

Woodford, M. (2012). Inflation Targeting and Financial Stability. *National Bureau of Economic Research Working Paper Series*, No. 17967.

Zimmermann, K. (2019). Monetary policy and bank profitability, 1870 – 2015. *Macroeconomics: Monetary & Fiscal Policies eJournal*.

Appendix A

Supplemental Material to Chapter 1

A.1 Existing Theoretical Models

This appendix briefly summarises how some of the different theoretical models predict that a contractionary monetary policy shock decreases bank leverage.

1. Woodford (2012), building on Curdia and Woodford (2010), uses a New Keynesian model to document that accomodative monetary policy increases the financial institution leverage. In his model, this increases the probability of a crisis by assumption. The mechanism relies on a postulated law of motion whereby leverage depends positively on the output gap. Therefore, a contractionary shock would contract output and subsequently leverage.

2. Dell'Ariccia et al. (2014) develop a model of financial intermediation where banks can engage in costly monitoring to reduce the credit risk in their loan portfolios. Monitoring effort and the pricing of bank assets and liabilities are endogenously determined. In equilibrium, they depend on the risk-free real interest rate (i.e., policy rate). Banks have limited liability and so take excessive risk which induces investors to enforce a leverage requirement. If the policy rate increases, this raises the rate the bank pays on debt liabilities and so exacerbates the agency problem. Investors therefore require banks to have more 'skin-in-the-game' to reduce this moral hazard and enforce tighter leverage requirements. Thus, the model features higher interest rates inducing lower bank leverage.

3. Drechsler et al. (2018b) develop a dynamic asset pricing model in which monetary policy affects the risk premium component of the cost of capital. Risk-tolerant agents (banks) borrow from risk-averse agents (i.e., take deposits) to fund levered investments. Leverage exposes banks to funding shocks. As such, banks hold liquidity buffers (e.g., US Treasuries) to insure against such funding shocks. If the central bank raises interest rates, the cost of holding liquid securities increase (i.e., there is a higher liquidity premium). This increase in the price of funding shock insurance means banks will reduce their liquidity buffers. Therefore, with lower insurance, banks reduce their exposure to funding shocks by reducing leverage. Hence an increase in the central bank rate reduces bank leverage.

4. Martinez-Miera and Repullo (2019), building on Martinez-Miera and Repullo (2017), in which competitive financial institutions that are funded with uninsured debt can engage in costly monitoring of entrepreneurial firms. However, monitoring is unobservable, so there is a moral hazard problem. They also include the possibility of costly equity financing for banks where greater equity can ameliorate the moral hazard problem. They find that tightening monetary policy reduces the wealth that investors allocate to funding entrepreneurs and banks. This decreases aggregate investment and lowers the return on debt and equity and ultimately increases leverage.

A.2 Book Leverage versus Market Leverage

In this paper, I use accounting-based measures of leverage (i.e., book leverage). An alternative approach would be to use market-based measures of leverage. Each measure has its own advantages and disadvantages. The definition of book leverage is the ratio of total assets to book equity while the definition of market leverage is the ratio of enterprise value (i.e., the sum of total liabilities and market equity) to market equity where market equity captures the market value of equity. I use book leverage for several reasons.

The first reason is consistency with the overall policy framework. When considering financial stability, macroprudential regulations focus on book leverage rather than market leverage. As such, from a policy consistency perspective, one would expect that monetary policy that targets financial stability would also do so through book leverage.

The second reason relates to bank decision-making. Banks themselves present their targets for return on equity at book value and report the evolution of leverage at book value. Indeed, Adrian et al. (2019) documents empirically that banks base their balance sheet management around book equity and book leverage and as such actively manage book leverage. While they mention market leverage also plays a role, they conclude that it is secondary to book leverage determined primarily by market forces. Similarly, Li (2022) highlight that it is book leverage that matters for bank lending decisions. Nuño and Thomas (2017) also highlight that book equity is the appropriate notion of equity when interested in the bank lending while market equity would be more appropriate if interested in new share issuance or mergers and acquisitions decisions. Given the role of book leverage in lending decisions, it clearly interacts more directly with the bank lending channel of monetary policy and would therefore constitute the appropriate measure of leverage for my analysis.

The third reason relates to explicit modelling choices. While many papers do not explicitly model book leverage or market leverage, they often implicitly consider book leverage. For example, models that rely on procyclicality of leverage are considering book leverage as market leverage is countercyclical. Ottonello and Song (2022) show analytically that in their model there is a tight link between book leverage and market leverage. More recently, Begenau and Landvoigt (2021) construct a rich model where delayed loss recognition can explain why book values differ from fundamental values.

The final reason is a question of data. Book leverage captures the entirety of the banking system as this data is available for all banks. Including the entire system is important in order to most accurately evaluate aggregate macroeconomic effects. Market leverage is only available for listed banks and so would significantly narrow the scope of the analysis.

Despite its limited scope, I repeat my analysis using market leverage. Figure A.6 shows that the results are qualitatively similar, albeit noisier and larger in magnitude for market leverage. The measure of market leverage I use is from He et al. (2017). They construct it as follows:

$$\text{Market Leverage}_t = \frac{\sum_i (\text{Market Equity}_{i,t} + \text{Book Debt}_{i,t})}{\sum_i \text{Market Equity}_{i,t}} \tag{A.1}$$

A few reasons for the different response in terms of magnitude are that the measure is only for the bank-holding companies of primary dealers. My data for book leverage is at the commercial bank level. Therefore, the samples are not strictly comparable. However, it is not especially surprising that they yield similar qualitative results as He et al. (2017) highlight that book and market leverage exhibit a strong positive correlation for the primary dealers in their sample.

A.3 Robustness Checks

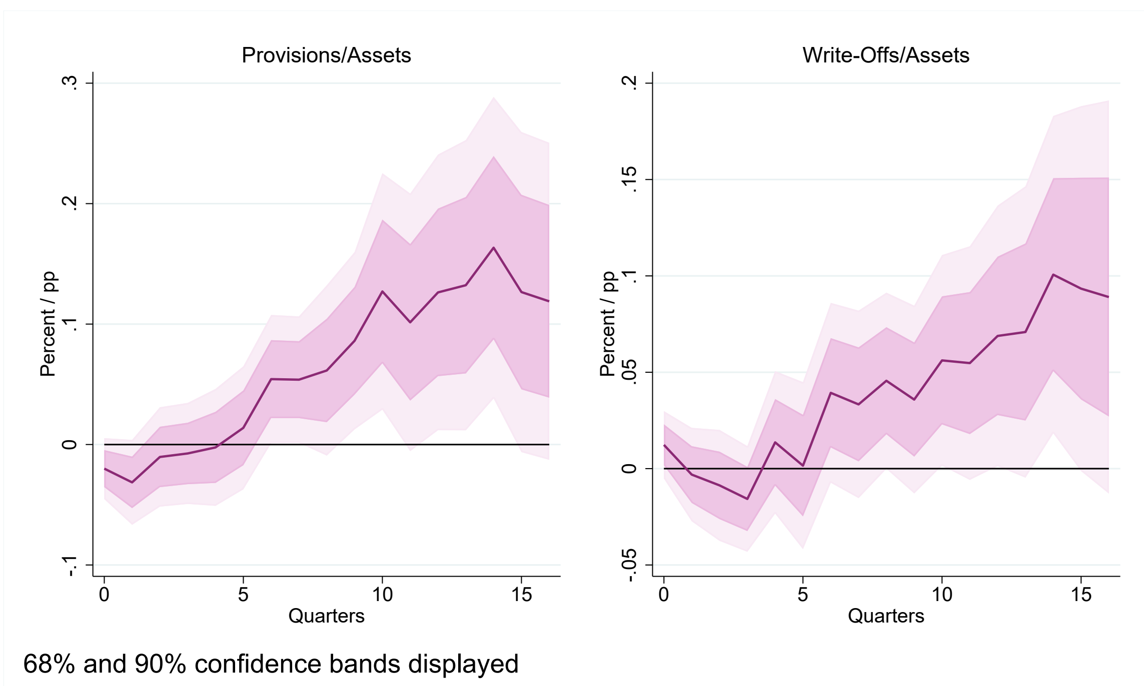

Figure A.1. Impulse Response of Provisions and Write-Offs

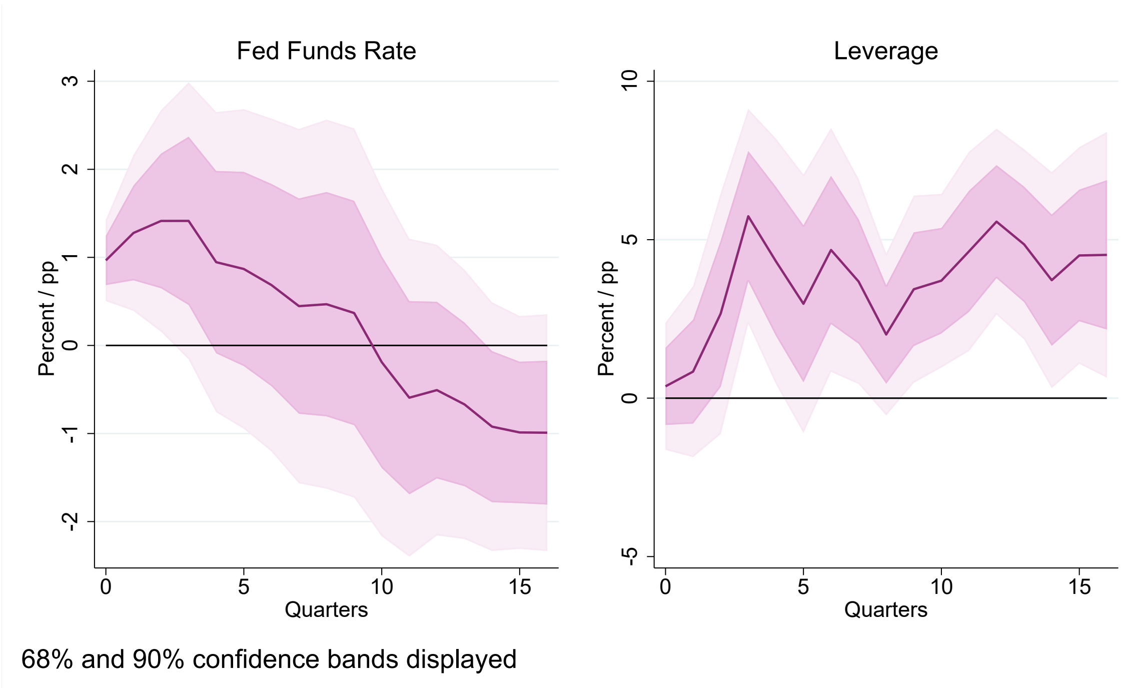

Figure A.2. Impulse Response of Regulatory Leverage to Contractionary Monetary Shock

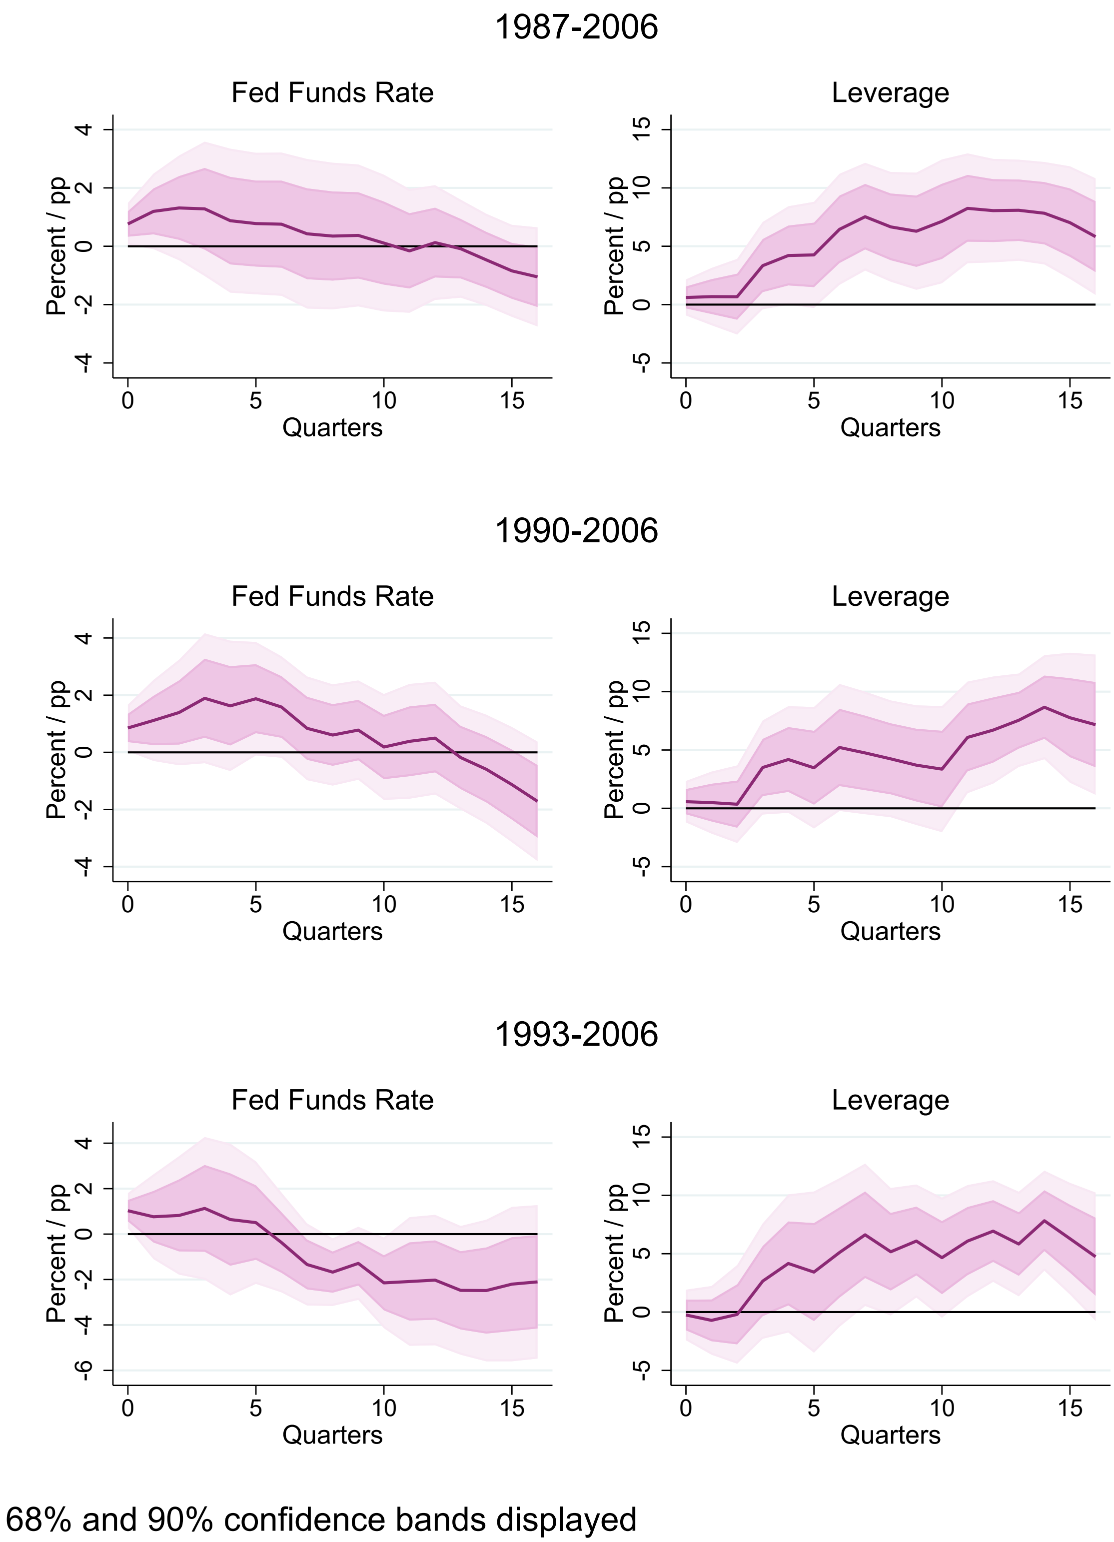

Figure A.3. Different Time Periods

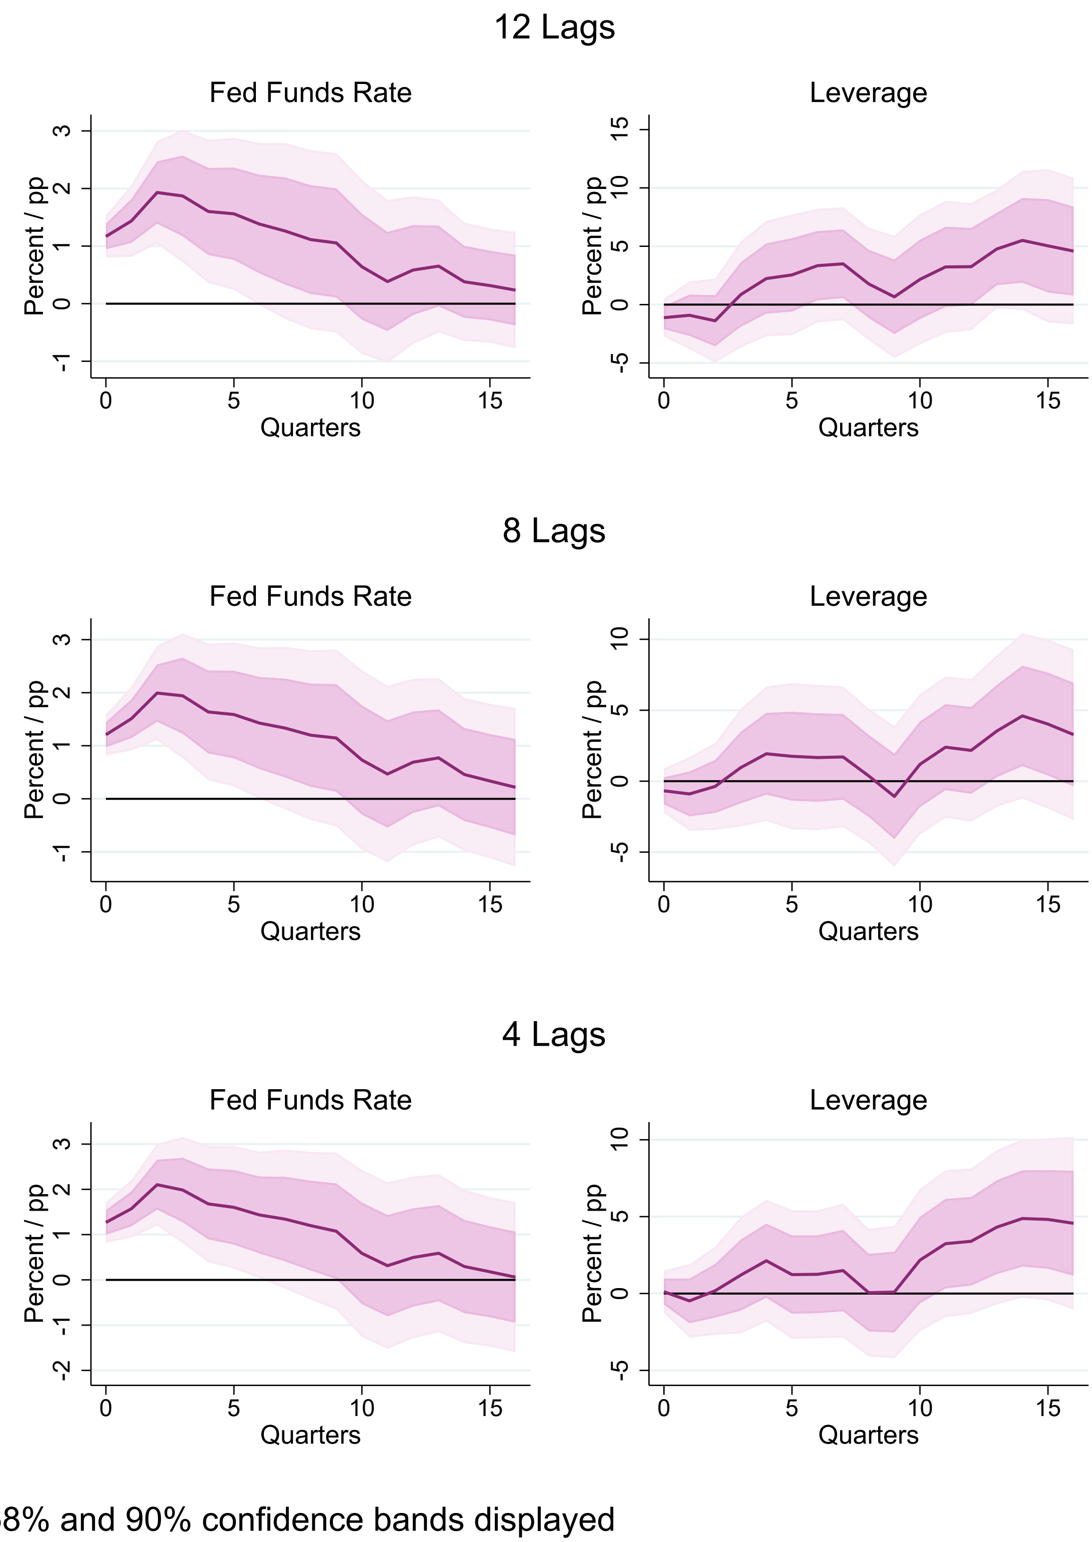

Figure A.4. Different Number of Lags

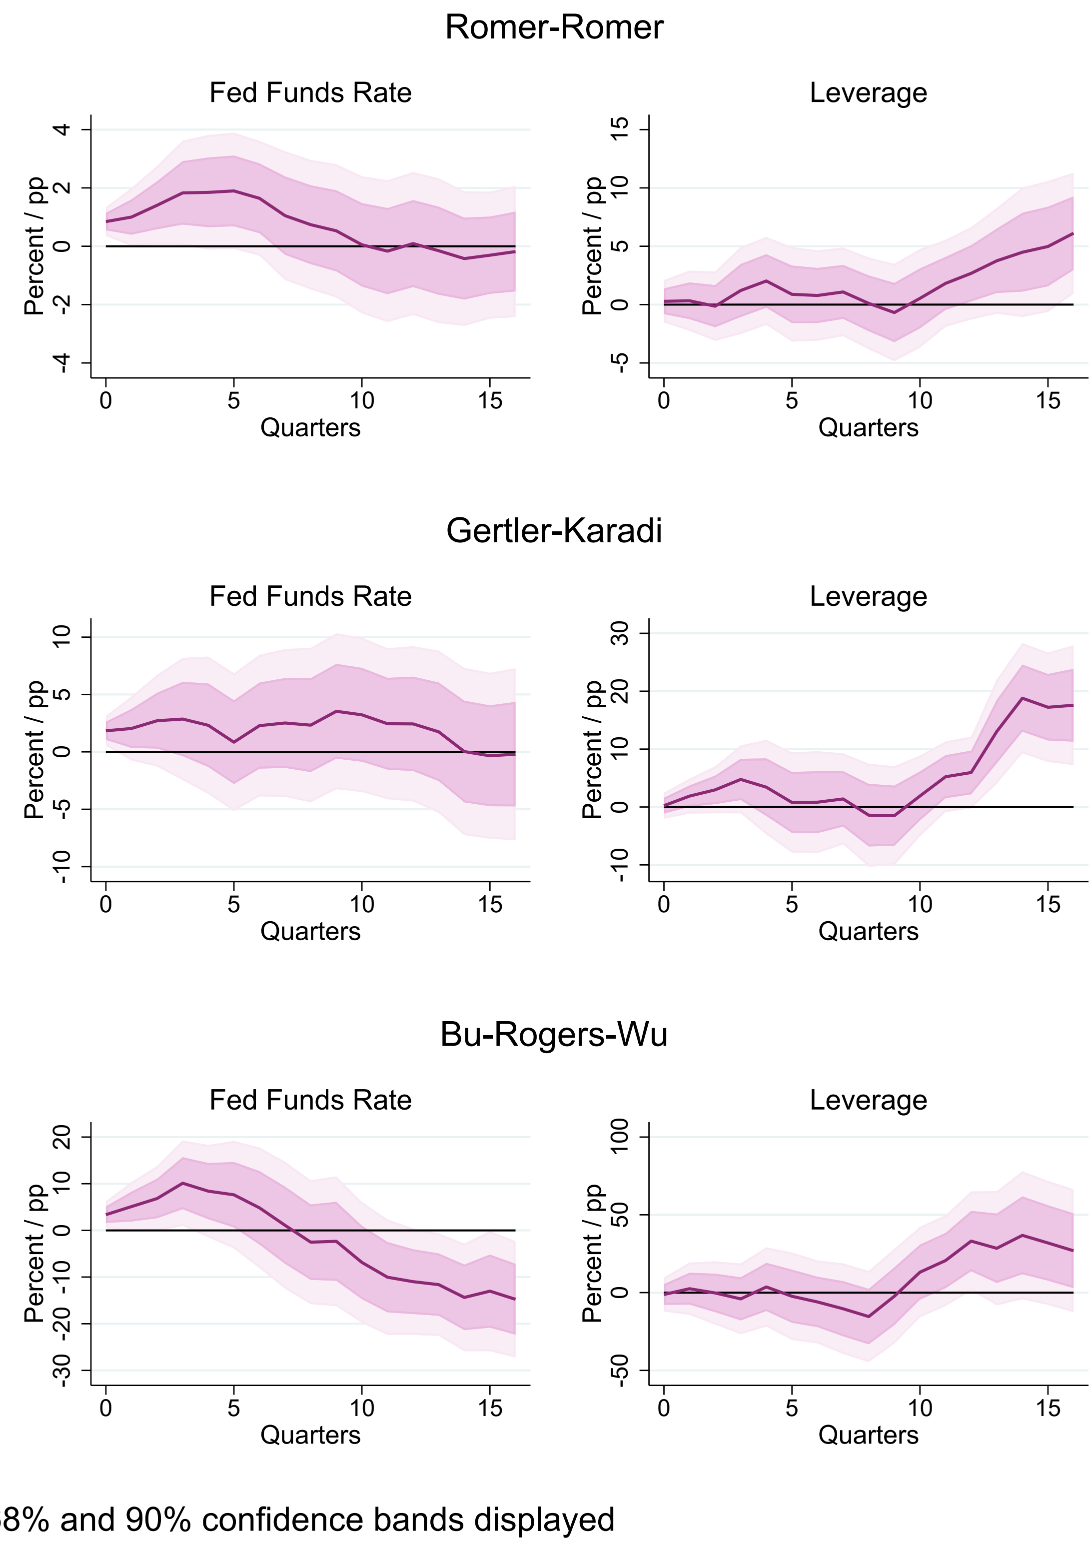

Figure A.5. Book Leverage Response to different Monetary Policy Shock Series

Figure A.6. Market Leverage Response to different Monetary Policy Shock Series

A.4 Theoretical Model

The model builds on Kirti (2020) but incorporates credit risk through loan losses on floating-rate loans.

A.4.1 Bank Problem

The bank has the following objective

$$\max_{f_L} V_b = E[\pi_b] - \frac{\gamma}{2} Var[\pi_b] \tag{A.2}$$

where profits are given by the following

$$\pi_b = L(1-f_L)(\bar{r}+\mu(f_L)) + Lf_L(\bar{r}+\varepsilon+\mu(f_L)) - D(\bar{r}+\beta\varepsilon) - Lf_L\theta(\varepsilon) \tag{A.3}$$

Therefore, we can rewrite V_b as

$$\begin{aligned} V_b = {} & L(1-f_L)\bar{r} + Lf_L\bar{r} + L\mu(f_L) - D\bar{r} - Lf_L\overline{\theta(\varepsilon)} - \frac{\gamma L^2 f_L^2 \sigma_\varepsilon^2}{2} - \frac{\gamma D^2 \beta^2 \sigma_\varepsilon^2}{2} - \frac{\gamma L^2 f_L^2 \sigma_\theta^2}{2} \\ & + \gamma L f_L D \beta \sigma_\varepsilon^2 + \gamma L^2 f_L^2 \rho_{\varepsilon\theta} - \gamma D \beta L f_L \rho_{\varepsilon\theta} \end{aligned} \tag{A.4}$$

where $\sigma_\varepsilon^2 = Var[\varepsilon]$, $\sigma_\theta^2 = Var[\theta(\varepsilon)]$, $E[\theta(\varepsilon)] = \overline{\theta(\varepsilon)}$, and $Cov(\varepsilon, \theta(\varepsilon)) = \rho_{\varepsilon\theta}$. Note that I assume the following: $\sigma_\varepsilon^2 > \sigma_\theta^2$ and $\rho_{\varepsilon\theta} > 0$ where the latter captures that the loan-loss rate increases in the size of the monetary policy shock.

Taking the first-order condition with respect to f_L and simplifying yields the following expression for f_L^*

$$f_L^* = \frac{\frac{\partial \mu(f_L)}{\partial f_L} - \overline{\theta(\varepsilon)}}{\gamma L\left(\sigma_\varepsilon^2 + \sigma_\theta^2 - 2\rho_{\varepsilon\theta}\right)} + \frac{D\beta\left(\sigma_\varepsilon^2 - \rho_{\varepsilon\theta}\right)}{L\left(\sigma_\varepsilon^2 + \sigma_\theta^2 - 2\rho_{\varepsilon\theta}\right)} \tag{A.5}$$

Note that the term in parentheses in the denominator is positive as it is simply the variance of the difference between the monetary policy shock and the loan-loss rate. Therefore, the denominator is also positive. Moreover, the numerator in the second term is positive as $\sigma_\varepsilon^2 + \sigma_\theta^2 - 2\rho_{\varepsilon\theta} > 0$ and $\sigma_\varepsilon^2 > \sigma_\theta^2$, so $\sigma_\varepsilon^2 > \rho_{\varepsilon\theta}$.

All else equal, a bank would choose a higher floating share if it is more exposed to interest expense on its deposits (e.g., through a higher deposit-loan ratio or a higher deposit beta). This is because the floating share would act as a hedge. However, the bank will choose a lower floating share if it more exposed to credit risk from monetary policy shocks (e.g., through a higher $\overline{\theta(\varepsilon)}$). This is because the hedge comes at the cost of credit risk. The specific functional form of $\theta(\varepsilon)$ and its covariance with the shock will determine the sensitivity of these effects.

A.4.2 Firm Problem

The firm has a similar objective function (with the same risk-aversion coefficient), except that it is choosing how much invest, I, which it can only do through borrowing. So the firm objective function is

$$\max_I V_f = E[\pi_f] - \frac{\gamma}{2} Var[\pi_f] \tag{A.6}$$

where firm profits are given by the following

$$\pi_f = AI - I - I(1 - f_L)(\bar{r} + \mu(f_L)) - If_L(\bar{r} + \varepsilon + \mu(f_L)) - If_L\theta(\varepsilon) \tag{A.7}$$

Note that $If_L\theta(\varepsilon)$ captures in, a reduced form way, that the firm cannot repay some of its floating-rate debt if there is a contractionary monetary policy shock.

We can now rewrite V_f as the following

$$V_f = AI - I - I\bar{r} - I\mu(f_L) - If_L\overline{\theta(\varepsilon)} - \frac{\gamma}{2}\left(I^2 f_L^2 \sigma_\varepsilon^2 + I^2 f_L^2 \sigma_\theta^2 + 2If_L\rho_{\varepsilon\theta}\right) \tag{A.8}$$

Taking the first-order condition with respect to I and simplifying yields the following expression for $\mu(f_L)$

$$\mu(f_L) = A - 1 - \bar{r} - f_L\overline{\theta(\varepsilon)} - \gamma I f_L^2 \sigma_\varepsilon^2 - \gamma I f_L^2 \sigma_\theta^2 - \gamma f_L \rho_{\varepsilon\theta} \tag{A.9}$$

A.4.3 Equilibrium

In equilibrium, we will have a loan spread, μ^* that will equate firm credit demand, I, with bank loan size, L. So, using $I = L$ and plugging the derivative of (A.9) with respect to f_L into (A.5) yields the equilibrium f_L^*

$$f_L^* = \frac{D\beta\gamma\left(\sigma_\varepsilon^2 - \rho_{\varepsilon\theta}\right) - \gamma\rho_{\varepsilon\theta} - 2\overline{\theta(\varepsilon)}}{\gamma L\left(3\sigma_\varepsilon^2 + 3\sigma_\theta^2 - 2\rho_{\varepsilon\theta}\right)} \tag{A.10}$$

Note that the denominator of (A.10) is positive because γ, L, and $3\sigma_\varepsilon^2 + 3\sigma_\theta^2 - 2\rho_{\varepsilon\theta}$ are all positive.[1] Moreover, given that D and β are positive, and that $\sigma_\varepsilon^2 > \rho_{\varepsilon\theta}$, then we have $\frac{\partial f_L}{\partial \beta} > 0$, which is consistent with banks using floating-rate loans as a hedge against interest rate risk. One can also show that f_L^* is positive as $\frac{\rho_{\varepsilon\theta}}{D\beta} \approx 0$, $\frac{2\overline{\theta(\varepsilon)}}{D\beta\gamma} \approx 0$, and $\sigma_\varepsilon^2 > \rho_{\varepsilon\theta}$.[2]

[1] $3\sigma_\varepsilon^2 + 3\sigma_\theta^2 - 2\rho_{\varepsilon\theta}$ is positive because $3\sigma_\varepsilon^2 + 3\sigma_\theta^2 - 2\rho_{\varepsilon\theta} > \sigma_\varepsilon^2 + \sigma_\theta^2 - 2\rho_{\varepsilon\theta} \equiv Var(\varepsilon - \theta) > 0$

[2] Note this requires that σ_ε^2 has to be sufficiently large relative to $\rho_{\varepsilon\theta}$.

Appendix B

Supplemental Material to Chapter 2

B.1 Additional Tables

Table B.1. Daily change in Treasury yield, 1994-2016

	1-year Treasuries	2-year Treasuries	10-year Treasuries
Dummy=1 in Week 0	-0.000672	0.0000238	-0.00362
	(-0.39)	(0.01)	(-1.56)
Dummy=1 in Week 2,4,6	0.000127	-0.00131	-0.00199
	(0.10)	(-0.83)	(-1.18)
N	5997	5997	5997

t statistics in parentheses
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Table B.2. Daily change in futures yield, 1994-2016

	2nd FF Contract	4th FF Contract	4th ED contract
Dummy=1 in Week 0	-0.000349	-0.000404	-0.00301
	(-0.39)	(-0.36)	(-1.04)
Dummy=1 in Week 2,4,6	-0.000298	-0.000245	-0.00334
	(-0.49)	(-0.31)	(-1.60)
N	5996	5996	5996

t statistics in parentheses
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Table B.3. Expanding Window Regressions 1994 - 2023

Time Period	Even-Week Regression Coefficient	*t* statistic	N
1994 - 2023	0.0801***	3.02	7821
1995 - 2023	0.0790***	2.89	7561
1996 - 2023	0.0806***	2.86	7301
1997 - 2023	0.0848***	2.91	7039
1998 - 2023	0.0830***	2.78	6778
1999 - 2023	0.0782**	2.57	6517
2000 - 2023	0.0708**	2.27	6256
2001 - 2023	0.0593*	1.88	5996
2002 - 2023	0.0685**	2.14	5735
2003 - 2023	0.0708**	2.19	5474
2004 - 2023	0.0557*	1.67	5213
2005 - 2023	0.0592*	1.71	4952
2006 - 2023	0.0571	1.57	4694
2007 - 2023	0.0671*	1.76	4434
2008 - 2023	0.0689*	1.73	4173
2009 - 2023	0.0401	1.09	3911
2010 - 2023	0.0318	0.86	3650
2011 - 2023	0.0177	0.47	3389
2012 - 2023	0.0056	0.15	3129
2013 - 2023	0.0097	0.24	2868
2014 - 2023	0.0174	0.40	2607
2015 - 2023	0.0105	0.22	2346
2016 - 2023	-0.0200	-0.39	2085
2017 - 2023	-0.0504	-0.88	1824
2018 - 2023	-0.0710	-0.93	1565
2019 - 2023	-0.0277	-0.37	1304
2020 - 2023	-0.0224	-0.25	1043
2021 - 2023	0.0100	0.12	781
2022 - 2023	0.1120	1.01	520

* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Each line represents a separate regression with the same variables. The only change is the time period. The results in Table 2.3 differ from those here as the former separates out the week 0 dummy.

Table B.4. Day of Cycle Regression, 1994-2003

FOMC Cycle Day	Day of Cycle Coefficient	*t* statistic	N
-6	0.011	0.09	2608
-5	-0.137	-1.19	2608
-4	-0.115	-1.05	2608
-3	0.024	0.17	2608
-2	-0.053	-0.52	2608
-1	0.159	1.33	2608
0	0.191*	1.75	2608
1	0.243*	1.75	2608
2	-0.037	-0.35	2608
3	-0.097	-0.81	2608
4	-0.065	-0.54	2608
5	-0.142	-1.18	2608
6	-0.104	-0.88	2608
7	-0.075	-0.61	2608
8	0.076	0.63	2608
9	-0.316**	-2.49	2608
10	-0.044	-0.30	2608
11	0.070	0.55	2608
12	0.125	0.92	2608
13	0.263**	2.05	2608
14	-0.019	-0.16	2608
15	0.049	0.34	2608
16	-0.252**	-2.22	2608
17	0.018	0.13	2608
18	-0.079	-0.53	2608
19	-0.205	-1.12	2608
20	0.449***	2.95	2608
21	0.058	0.51	2608
22	0.054	0.35	2608
23	0.115	0.96	2608
24	-0.214	-1.52	2608
25	0.203	1.27	2608
26	0.059	0.38	2608
27	-0.368**	-2.20	2608
28	-0.074	-0.40	2608
29	0.363*	1.78	2608
30	-0.176	-0.63	2608
31	-0.101	-0.42	2608
32	0.131	0.53	2608
33	0.614*	1.80	2608

* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Each line represents a separate regression where the *x* variable is a different FOMC cycle day.

B.2 Additional Figures

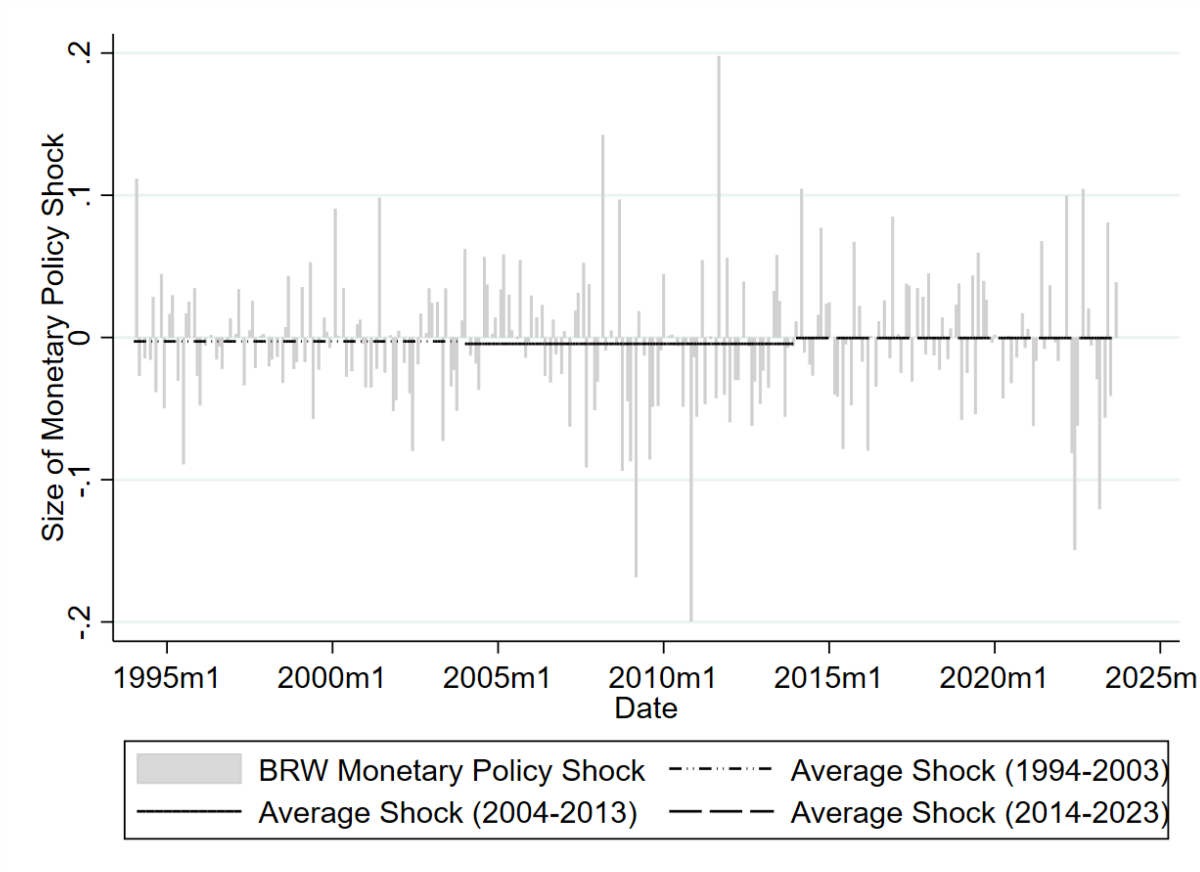

Figure B.1. Bu et al. (2021) Monetary Policy Shock, 1994-2023

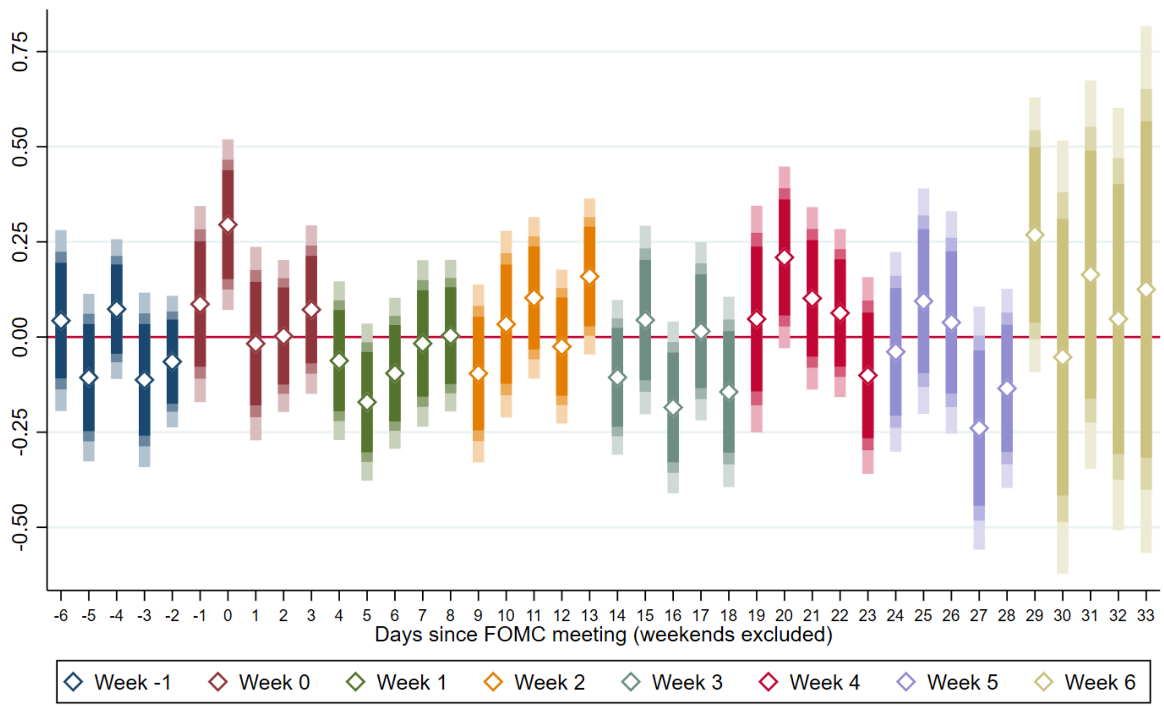

Figure B.2. Regressing Daily Excess Return on Day of the Cycle, 1994-2016

Notes: The 40 regressions underpinning this chart have excess returns as the dependent variable. The only regressors are a constant and the dummy for the specific day of the cycle. The diamonds reflect the coefficient on the day dummy of each regression. The confidence intervals are shown by coloured bars. The darkest shade represents the 99% confidence interval, one shade lighter represents the 95% confidence interval, and the lightest colour represents the 90% confidence interval.

Appendix C

Supplemental Material to Chapter 3

C.1 Proofs

Proof of proposition 1

Assume that a PBE of the game exists and let $\sigma = (\sigma_m, \sigma_x, \sigma_r)$ be the corresponding strategy profile, consisting of a communication rule $\sigma_m : [-\phi_1, \phi_1] \to \Delta(M)$, a rate rule $[-\phi_1, \phi_1] \times [-\phi_1, \phi_1] \times X^I \to \Delta(\mathbb{R})$ for *CB*, and an investment plan $x_i : M \to \Delta(X)$ for each $i \in I$.[1] Fix any $(\omega_1, \omega_2, m, x) \in [-\phi_1, \phi_1] \times [-\phi_2, \phi_2] \times M \times X^I$, and consider the problem of the *CB* at $t = 2$ in the subgame identified by $(\omega_1, \omega_2, m, x)$. In any PBE, the *CB* must pick $\hat{r}$ solving

$$\min_{r \in \mathbb{R}} L_{CB}(x, r, \omega)$$

for $\omega = \omega_1 + \omega_2$. Taking the derivative of the objective function with respect to r one obtains

$$\frac{\partial L_{CB}}{\partial r} = (1-\alpha)(r-\omega) + \alpha g'(\bar{S}(x,r))(r-\bar{x})$$

for $\bar{x}$ a weighed average of $x_1, \dots, x_I$. Given assumption 1, the first order condition is necessary and sufficient for optimality, and gives the unique PBE rate choice $\sigma_r^\star(\omega_1, \omega_2, x)$ made by *CB* in any subgame $(\omega_1, \omega_2, m, x)$ reached at $t = 2$ as a function of the corresponding (ω_1, ω_2, x). From the first order condition, $\sigma_r^\star(\omega_1, \omega_2, x)$ must satisfy

[1] We only look at rate rules which are not contingent on messages $m \in M$, as doing otherwise is not meaningful in this static setting.

$$\sigma_r^\star(\omega_1,\omega_2,x) = \frac{1-\alpha}{(1-\alpha)+\alpha g'(\bar{S}(x,r))}\omega + \frac{\alpha g'(\bar{S}(x,r))}{(1-\alpha)+\alpha g'(\bar{S}(x,r))}\bar{x}. \tag{C.1}$$

Since this must hold in any PBE, it has $\sigma_r = \sigma_r^\star$. Note that communication and investment decisions are not functions of ω_2, since they are taken before ω_2 is realized. Hence for each $(\omega_1,\omega_2) \in [-\phi_1,\phi_1] \times [-\phi_2,\phi_2]$ and $m \in \text{supp}(\sigma_m(\omega_1))$, and $x_i \in \text{supp}(\sigma_{x_i}(m))$,

$$\frac{\partial \sigma_r^\star(\omega_1,\omega_2,x)}{\partial \omega_2} = -\frac{\frac{\partial^2 L_{CB}}{\partial \omega_2 \partial r}}{\frac{\partial^2 L_{CB}}{\partial r^2}} = \frac{1-\alpha}{(1-\alpha)+\alpha\left[2g''(\bar{S}(x,r))(r-\bar{x})^2 + g'(\bar{S}(x,r))\right]},$$

so that $0 < \frac{\partial \sigma_r^\star}{\partial \omega_2} < 1$ follows from $\alpha \in (0,1)$, $g' > 0$ and $g'' \geq 0$. Hence $0 < \frac{\partial r^\star(\sigma,\omega_1,\omega^2)}{\partial \omega_2} < 1$ for each $(\omega_1,\omega_2) \in [-\phi_1,\phi_1] \times [-\phi_2,\phi_2]$, proving proposition 1.

Proof of proposition 2

Let $\hat{\sigma} = (\hat{\sigma}_m, \hat{\sigma}_x, \hat{\sigma}_r) \in \hat{\Sigma}$. Let us show that conditions (i), (ii), and (iii) are necessarily satisfied by $\hat{\sigma}$ – starting from condition (iii) and proceeding backwards.

First, note that property (iii) is equivalent to requiring $\hat{\sigma}_r = \sigma_r^\star$ on path. Assume that there exist $(\omega_1,\omega_2,x) \in [-\phi_1,\phi_1] \times [-\phi_2,\phi_2] \times \mathbb{R}$ such that x arises with positive probability on path given $\hat{\sigma}$ and ω_1, and that $\hat{\sigma}_r(\omega_1,\omega_2,x) \neq \sigma_r^\star(\omega_1,\omega_2,x)$. Let $\sigma' = (\hat{\sigma}_m, \hat{\sigma}_x, \sigma_r^\star)$. By definition of $\sigma_r^\star$ and equation C.1, $L_{CB}(x', \sigma_r^\star(\omega_1', \omega_2', x'), \omega') \geq L_{CB}(x', \hat{\sigma}_r(\omega_1', \omega_2', x'), \omega')$ for each $(\omega_1', \omega_2', x') \in [-\phi_1,\phi_1] \times [-\phi_2,\phi_2] \times \mathbb{R}$, with $L_{CB}(x, \sigma_r^\star(\omega_1,\omega_2,x), \omega) > L_{CB}(x, \hat{\sigma}_r(\omega_1,\omega_2,x), \omega)$ by construction. Since (ω_1,ω_2,x) arises with positive probability given $(\hat{\sigma}_m, \hat{\sigma}_x)$ it must be that $W(\sigma') > W(\hat{\sigma})$, which contradicts $\hat{\sigma} \in \hat{\Sigma}$. This proves (iii) by contradiction.

For (ii), let $m \in \text{supp}\hat{\sigma}_m(\omega_1)$ for some $\omega_1 \in [-\phi_1,\phi_1]$. Knowing from (iii) that $\hat{\sigma}_r = \sigma_r^\star$, and using $g'' = 0$ we have that

$$\begin{aligned}\mathbb{E}[L_{CB}(x, \sigma_r^\star(\omega_1,\omega_2,x), \omega)|m, \hat{\sigma}_m] = \frac{1}{2}(1-\alpha)\left[\frac{\alpha k(x)}{(1-\alpha)+\alpha k(x)}\right]^2 \mathbb{E}[(\omega-\bar{x})^2|m,\hat{\sigma}_m] \\ + \frac{1}{2}\alpha\mathbb{E}\left[g(\bar{S}(x, \sigma_r^\star(\omega_1,\omega_2,x))|m,\hat{\sigma}_m\right]\end{aligned}$$

where $k: X \to \mathbb{R}$, $k(x) = g'(\bar{S}(x,r))$. Note that the above expression is convex in x_i, and therefore its unique minimizer $x \in \mathbb{R}^N$ is obtained by imposing the FOC for each $i \in I$. This yields the following set of FOCs, for $w_i \in [0,1]$ being the weight assigned to i in computing aggregate readjustments,

$$(1-\alpha)\left[\frac{\alpha k(x)}{(1-\alpha)+\alpha k(x)}\right]^2(\mathbb{E}[\omega_1|m,\hat{\sigma}_m]-\bar{x}) = \alpha k(x)\left[x_i - \frac{1-\alpha}{(1-\alpha)+\alpha k(x)}\mathbb{E}[\omega_1|m,\hat{\sigma}_m] - \frac{\alpha k(x)}{(1-\alpha)+\alpha k(x)}\bar{x}\right]\left(1-\frac{\alpha k(x)}{(1-\alpha)+\alpha k(x)}w_i\right), \quad \forall i \in I$$

which is solved by $x_i = \mathbb{E}[\omega_1|m,\hat{\sigma}_m]$ for each $i \in I$. But this means that $\sigma_{x_i}(m) = \mathbb{E}[\omega_1|m,\hat{\sigma}_m]$ must be true of all profiles in $\hat{\Sigma}$, regardless of $\hat{\sigma}_x$. This proves (ii).

Finally, to see why it must be that $\hat{\sigma}_m$ fully reveals ω_1, note for $\hat{\sigma}_x, \hat{\sigma}_r$ satisfying properties (ii) and (iii), and given $g'' = 0$, $\mathbb{E}[L_{CB}|\sigma_m\hat{\sigma}_x,\hat{\sigma}_r]$ is increasing in the residual variance $\hat{\sigma}_1^2$ of ω_1 induced by the communication strategy σ_m. Hence it must be that $\hat{\sigma}_m$ is fully revealing.

To prove that in the game described in section 3 properties (i), (ii), (iii) are sufficient for a strategy profile to belong to $\hat{\Sigma}$, it is sufficient to notice that any profile satisfying the three properties leads to the same ex-ante payoff $\hat{W} = -\frac{1}{2}\alpha(1-\alpha)\sigma_2$.

The efficiency of the competitive equilibrium is obtained by letting $N \to \infty$ in the equilibrium strategy profile σ^{oli} presented in proposition 3 and proposition 4. Specifically, note that $\sigma_r^{oli} = \sigma_r^*$ so that (iii) is satisfied. The bias term in proposition 3.i goes to 0 as $N \to \infty$ proving that unbiasedness of $\sigma_{x_i}^{oli}$ holds to the limit, satisfying (ii). Finally note that the expression for $\bar{P}$ in proposition 4.ii goes to ∞ as $N \to \infty$, making equilibrium communication σ_m^{oli} fully informative, and satisfying (i) to the limit.

Proof of proposition 3

We start from part (i). Fix a communication rule $\sigma_m \in \Sigma_m$. We now derive the oligopolistic best responses $\sigma_{x_i}^{oli}$ and σ_r^{oli}, where σ_r^{oli} is a best response to $(\sigma_x^{oli}, \sigma_m)$ and, for each $i \in I$, $\sigma_{x_i}^{oli}$ is a best response to $(\sigma_m, \sigma_r^{oli})$.

First, note that it must be that $\sigma_r^{oli} = \sigma_r^\star$ holds pointwise, where $\sigma_r^\star$ satisfies equation C.1. Setting $\tilde{x}_i \equiv x_i$ in C.1 by definition, simple algebra yields

$$\sigma_r^\star(\omega_1, \omega_2, x) = (1-\alpha)\omega + \alpha \bar{x} \tag{C.2}$$

for each $(\omega_1, \omega_2, x) \in [-\phi_1, \phi_1] \times [-\phi_2, \phi_2] \times X^I$, where $\bar{x} = \frac{1}{N}\sum_{i \in I} x_i$.

Plugging C.2 in 3.1, the payoff of i as a function of $(\omega_1, \omega_2, x) \in [-\phi_1, \phi_1] \times [-\phi_2, \phi_2] \times X^I$ becomes

$$u_i(\sigma_r^\star(\omega_1, \omega_2, x), x_i) = -\frac{1}{2}[x_i - (1-\alpha)\omega - \alpha\bar{x}]^2 - \beta(1-\alpha)\omega - \alpha\beta\bar{x}.$$

At $t = 1$ and after any message $m \in M$, each i chooses a $\sigma_{x_i}^{oli}(m)$ supported on the set of maximizers of $\mathbb{E}[u_i(\sigma_r^\star(\omega_1, \omega_2, x), x_i)|m, \sigma_{x_{-i}}(m)]$ where the expectation is based on a posterior belief on ω_1 that satisfies Bayes rule whenever possible, and on the strategy profile $\sigma_{x_{-i}}$ played by all other investors. As the objective function is strictly concave, the unique minimizer $x_i^\star$ satisfies the FOC,

$$[x_i^\star - (1-\alpha)\mathbb{E}[\omega_1|m, \sigma_m] - \alpha\bar{x}^e]\left(1 - \frac{\alpha}{N}\right) + \frac{\alpha\beta}{N} = 0 \tag{C.3}$$

where $\bar{x}^e = \frac{x_i^\star}{N} + \sigma_{x_{-i}}(x'_{-i}|m)\sum_{j \in I\setminus\{i\}} \frac{x'_j}{N}$. First, note that the uniqueness of the minimizer implies that no mixed strategies are played. Second, the fact that $x_i^\star$ depends on x_j, $j \neq i$, only through the average investment $\bar{x}$ implies symmetry of investment whenever m sent with positive probability from σ_m.[2] Hence, we can let $\bar{x}^e = x_i^\star$ in the condition above to obtain that the optimal strategy of each $i \in I$ after message m occurring with positive probability given σ_m. In particular, the profile

[2]For m sent with positive probability form σ_m the expectation term in equation (C.3) takes the same value for all investors, but after a surprising m the expectation term might differ across investors since Bayes rule does not apply.

best responses σ_x^{oli} must satisfy

$$\sigma_{x_i}^{oli}(m) - \mathbb{E}[\omega_1 | m, \sigma_m] = -\frac{\alpha\beta}{(N-\alpha)(1-\alpha)} \tag{C.4}$$

with proves part (i).

For part (ii), let $\sigma_m = \sigma_m^{com}$ for some σ_m^{com} fully revealing the state and let σ_x^{oli} satisfy condition (C.4) with respect to σ_m^{com}. On the path of play we have that for each $\omega_1 \in [-\phi_1, \phi_1]$, $m \in \text{supp}\sigma_m^{tr}(\omega_1)$ and each $i \in I$, $\sigma_{x_i}^{oli}(m) = \omega_1 - \frac{\alpha\beta}{(N-\alpha)(1-\alpha)}$, so that

$$\bar{x} = \omega_1 - \frac{\alpha\beta}{(N-\alpha)(1-\alpha)}. \tag{C.5}$$

By plugging C.5 in C.2 we obtain $r^\star(\sigma^{oli,tr}, \omega_1, \omega_2)$, for $\sigma^{oli,tr} = (\sigma_m^{com}, \sigma_x^{oli}, \sigma_r^\star)$,

$$r^\star(\sigma^{oli,tr}, \omega_1, \omega_2) = \omega_1 + (1-\alpha)\omega_2 - \frac{\alpha^2\beta}{(N-\alpha)(1-\alpha)},$$

which proves part (ii).

Proof of proposition 4

First, note that in any PBE it must be that CB optimally plays the strategy given by (C.2) a $t = 2$. Fixing this rate rule, we can write the expected loss of CB given state $\omega_1 \in [-\phi_1, \phi_1]$ as a function of $x \in X^I$ is

$$\mathbb{E}[L_{CB}(x,\omega)|\omega_1] = \frac{1}{2}\alpha^2(1-\alpha)\mathbb{E}[(\bar{x}-\omega)^2|\omega_1] + \alpha\frac{1}{2N}\sum_{i\in I}\mathbb{E}\left[[(1-\alpha)\omega + \alpha\bar{x} - x_i]^2|\omega_1\right] \tag{C.6}$$

As shown in the previous proof, any equilibrium will induce $x_i = x_j$ for each $i, j \in I$ regardless of the communication strategy employed, so communication in this game cannot be used, in any an equilibrium, to induce heterogeneity in investment choices.[3] Therefore, we restrict the attention

[3] Per proposition 2.ii, CB would not benefit from heterogeneous investment plans $x \in X^I$, provided that plans are unbiased.

the profiles in the set $\bar{X} = \left\{x \in X : x_i = x_j \quad \forall i, j \in I\right\}$. For $x \in \bar{X}$, (C.6) simplifies to

$$\begin{aligned}\mathbb{E}[L_{CB}(x,\omega)|\omega_1] &= \frac{1}{2}\alpha(1-\alpha)\mathbb{E}[(\bar{x}-\omega)^2|\omega_1] \\ &= \frac{1}{2}\alpha(1-\alpha)(\bar{x}-\omega_1)^2 + \frac{1}{2}\alpha(1-\alpha)\sigma_2^2.\end{aligned}$$

It is immediate to see from the above expression, that, in any PBE, the *CB* problem simplifies to choosing some $\sigma_m^{oli} \in \Sigma_m$ that minimizes $\mathbb{E}[(\bar{x}-\omega_1)^2]$ *given* the PBE strategies $(\sigma_x^{oli}, \sigma_r^\star)$. But we know from (C.4) that, for each $i \in I$,

$$\sigma_{x_i}^{oli}(m) = \mathbb{E}[\omega_1|m, \sigma_m^{oli}] - \frac{\alpha\beta}{(N-\alpha)(1-\alpha)}$$

must be satisfied in equilibrium. Focusing on communication strategies σ_m for which each message occurs with positive probability in at least one state, the set of PBE communication strategies σ_m^{oli} is the argmax of the following problem

$$\begin{aligned}\min_{\sigma_m \in \Sigma_m} \int_{-\phi_1}^{\phi_1}\int_{-\phi_1}^{\phi_1} &(\sigma_{x_i}^{oli}(m)-\omega_1)^2\sigma_m(m|\omega_1)(2\phi_1)^{-1}dmd\omega_1 \\ \text{s.t.} \quad &\text{(i) } \sigma_{x_i}^{oli}(m) = \mathbb{E}[\omega_1|m,\sigma_m] - \frac{\alpha\beta}{(N-\alpha)(1-\alpha)}, \\ &\text{(ii) } \forall m \in M, \exists \omega_1 \in [-\phi_1, \phi_1] : \sigma_m(m|\omega_1) > 0.\end{aligned} \tag{C.7}$$

The general solution of this type of problem has been characterised by Crawford and Sobel (1982), hereafter referred to as CS. Note that relaxing (ii) would not deliver economically different communication strategies/equilibria provided that after messages m that are unexpected (i.e., with $\sigma_m(m|\omega_1) = 0$ for each $\omega_1 \in [-\phi_1, \phi_1]$), we restrict off-path beliefs to be such that for each $i \in I$ the best response to m does not expand the set of x_i played on the equilibrium path.

The mapping to CS is particularly evident by noticing that we can transform the variables and parameters of interest in our setting to match a specific case of their formulation. First, define the following transformation: $t : \mathbb{R} \to \mathbb{R}, t(x) = (2\phi_1)^{-1}\left(x + \phi_1 + \frac{\alpha\beta}{(N-\alpha)(1-\alpha)}\right)$. For each

$\omega_1 \in [-\phi_1, \phi_1]$ we can apply the transformation t to the perfect information ideal investment levels of CB and the representative investor i, denoted as $x^{\star}_{CB}(\omega_1) = \omega_1$ and $x^{\star}_i(\omega_1) = \omega_1 - \frac{\alpha\beta}{(N-\alpha)(1-\alpha)}$ respectively. It is easy to see that $t(x^{\star}_i(\omega_1)) = (2\phi_1)^{-1}(\omega_1+\phi_1)$ and $t(x^{\star}_{CB}(\omega_1)) = t(x^{\star}_i(\omega_1))+b$ for $b = \frac{\alpha\beta}{2\phi_1(N-\alpha)(1-\alpha)}$. Denote by $t^{\star}_{\omega_1}$ the random variable equal to $(2\phi_1)^{-1}(\omega_1+\phi_1)$, the transformed ideal point of i. First, note that $t^{\star}_{\omega_1} \sim U[0,1]$, supported on the unit interval as in CS. Second, note that for each $\omega_1 \in [-\phi_1, \phi_1]$ and $x_i \in \mathbb{R}$ it has

$$\left(\omega_1 - \frac{\alpha\beta}{(N-\alpha)(1-\alpha)} - x_i\right) \propto (t^{\star}_{\omega_1} - t(x_i))$$

and, similarly,

$$(\omega_1 - x_i) \propto (t_{\omega_1} + b - t(x_i))$$

Third, note that t is a bijection between X and X and is therefore invertible.

Substantially, t creates a“bridge” between our problem (C.7) and the seminal example studied by CS (section 4). Consider that example, and set $b = \frac{\alpha\beta}{2\phi_1(N-\alpha)(1-\alpha)}$. The above relations imply that (i) for every partition PBE of Example 4 where communication partitions $[0,1]$ in P intervals with cutoffs $a_0 < a_1 < ... < a_P$, there exists a PBE of our our game where σ^{oli}_m partitions $[-\phi_1, \phi_1]$ in P intervals, with cutoffs $t^{-1}(a_0+b) < t^{-1}(a_1+b) < ... < t^{-1}(a_P+b)$; and (ii) for every partition PBE of our game where σ^{oli}_m partitions $[-\phi_1, \phi_1]$ in P intervals with cutoffs $a_0 < a_1 < ... < a_P$, there exists a PBE of CS's seminal example, where communication partitions $[0,1]$ in P intervals, with cutoffs $t(a_0)-b < t(a_1)-b < ... < t(a_P)-b$. This correspondence between equilibrium communication strategies in the two games also implies that the strong comparative statics in section 4 of CS extend to our setting for $b = \frac{\alpha\beta}{2\phi_1(N-\alpha)(1-\alpha)}$, proving proposition 4.i-ii.

Proof of proposition 5

Let $b : [0,1) \to \mathbb{R}, b(\tilde{\alpha}) = -\frac{\tilde{\alpha}\beta}{(N-\tilde{\alpha})(1-\tilde{\alpha})}$, mapping central banker types to the corresponding investment bias in equilibrium (we ignore the case in which $\tilde{\alpha} = 1$ as such case leads to an infinite welfare loss). By repeating the steps of the proof of proposition 3 replacing α with $\tilde{\alpha} \in [0,1)$, it is easily shown that, for each $m \in M$, $(\omega_1, \omega_2) \in [-\phi_1, \phi_1] \times [-\phi_2, \phi_2]$ and $x \in X^I$,

$$\sigma_r^{\star}(\omega_1, \omega_2, x; \tilde{\alpha}) = (1-\tilde{\alpha})\omega + \tilde{\alpha}\bar{x}$$
$$\sigma_{x_i}^{oli}(m; \tilde{\alpha}) = \mathbb{E}[\omega_1 | m, \sigma_m] + b(\tilde{\alpha}).$$

It follows that the ex-ante loss for the central bank takes the following form, for $\sigma^{oli,tr}$ and $\sigma_{\bar{P}}^{oli}$ respectively,

$$W(\sigma^{oli,tr}(\tilde{\alpha})) = -\frac{1}{2}[\tilde{\alpha}^2(1-\alpha) + \alpha(1-\tilde{\alpha})^2](\sigma_2^2 + b(\tilde{\alpha})^2) \tag{C.8}$$
$$W(\sigma_{\bar{P}}^{oli}(\tilde{\alpha})) = -\frac{1}{2}[\tilde{\alpha}^2(1-\alpha) + \alpha(1-\tilde{\alpha})^2](\hat{\sigma}_{1,\bar{P}}^2(\tilde{\alpha}) + \sigma_2^2 + b(\tilde{\alpha})^2) \tag{C.9}$$

where the only difference is that in (C.8) the residual variance after communication is zero because communication is fully informative.

First, consider part (i) of the proposition. It is sufficient to compare $W(\sigma^{oli,tr}(\alpha))$ and $W(\sigma_{\bar{P}}^{oli}(\alpha))$ with $W(\sigma^{oli,tr}(0))$ and $W(\sigma_{\bar{P}}^{oli}(0))$ respectively, by plugging the corresponding values of $\tilde{\alpha}$) in C.8 and C.9. It has

$$W(\sigma^{oli,tr}(0)) > W(\sigma^{oli,tr}(\alpha)) \iff \sigma_2 < \sqrt{\frac{\alpha}{1-\alpha}}\left(\frac{\beta}{N-\alpha}\right) \equiv \sigma_2^{oli,tr}$$
$$W(\sigma^{oli}(0)) > W(\sigma^{oli}(\alpha)) \iff \sigma_2 < \sqrt{\frac{\alpha}{1-\alpha}\left(\frac{\beta}{N-\alpha}\right)^2 + \frac{1-\alpha}{\alpha}\hat{\sigma}_{1,\bar{P}}^2(\tilde{\alpha})} \equiv \sigma_2^{oli},$$

and indeed it holds $\sigma_2^{oli,tr} < \sigma_2^{oli}$ because $\hat{\sigma}_{1,\bar{P}}^2(\tilde{\alpha}) > 0$, which proves part (i).

Next, we prove part (ii). For the second statement, note that $W(\sigma^{oli,tr}(\tilde{\alpha})) \geq W(\sigma^{oli}(\tilde{\alpha}))$ for all $\tilde{\alpha} \in (0,1)$. Hence it must be $W(\sigma^{oli,tr}(\tilde{\alpha}^{oli})) \geq W(\sigma^{oli}(\tilde{\alpha}^{oli}))$. But by definition $W(\sigma^{oli,tr}(\tilde{\alpha}^{oli,tr})) \geq W(\sigma^{oli,tr}(\tilde{\alpha}^{oli}))$, implying $W(\sigma^{oli,tr}(\tilde{\alpha}^{oli,tr})) \geq W(\sigma^{oli}(\tilde{\alpha}^{oli}))$.

We now turn to the first statement of part (ii). First, let us take the derivative of (C.8) with respect to $\tilde{\alpha}$, yielding,

$$\frac{\partial W(\sigma^{oli,tr}(\tilde{\alpha}))}{\partial \tilde{\alpha}} = (\alpha-\tilde{\alpha})(\sigma_2^2+b(\tilde{\alpha})^2) - [\tilde{\alpha}^2(1-\alpha)+\alpha(1-\tilde{\alpha})^2]\underbrace{\frac{\partial b(\tilde{\alpha})}{\partial \tilde{\alpha}}b(\tilde{\alpha})}_{>0}.$$

On the one hand, the above expression is always negative for $\tilde{\alpha} \geq \alpha$, implying that the optimal central banker, if it exists, must be kitish. On the other hand, the derivative is positive for $\tilde{\alpha} = 0$, implying that $\tilde{\alpha}^{oli,tr} \in (0,\alpha)$, is the optimal central banker exists. Existence follows from the fact that $[0,\alpha]$ is compact and (C.8) continuous.

From the first order condition, we obtain following equality,

$$\sigma_2^2 + b(\tilde{\alpha})^2 = \left[\frac{\tilde{\alpha}^2(1-\alpha)+\alpha(1-\tilde{\alpha})^2}{\alpha-\tilde{\alpha}}\right]\frac{\partial b(\tilde{\alpha})}{\partial \tilde{\alpha}}b(\tilde{\alpha}) \tag{C.10}$$

Note that

$$\begin{aligned}
\frac{\partial b(\tilde{\alpha})}{\partial \tilde{\alpha}} &= -\frac{\beta(N-\tilde{\alpha})(1-\tilde{\alpha})+\tilde{\alpha}\beta(1+N-2\tilde{\alpha})}{(1-\tilde{\alpha})^2(N-\tilde{\alpha})^2} = -\frac{(N-\tilde{\alpha}^2)\beta}{(1-\tilde{\alpha})^2(N-\tilde{\alpha})^2}\\
\frac{\partial^2 b(\tilde{\alpha})}{\partial \tilde{\alpha}^2} &= 2\frac{\tilde{\alpha}\beta(1-\tilde{\alpha})(N-\tilde{\alpha})-(N-\tilde{\alpha}^2)\beta(1+N-2\tilde{\alpha})}{(1-\tilde{\alpha})^3(N-\tilde{\alpha})^3}\\
&= 2\frac{\tilde{\alpha}\beta N-\tilde{\alpha}^2\beta N-\tilde{\alpha}^2\beta+\tilde{\alpha}^3\beta-N\beta+\tilde{\alpha}^2\beta-N^2\beta+\tilde{\alpha}^2\beta N+2\tilde{\alpha}\beta N-2\tilde{\alpha}^3\beta}{(1-\tilde{\alpha})^3(N-\tilde{\alpha})^3}\\
&= 2\beta\frac{3\tilde{\alpha}N-N-N^2-\tilde{\alpha}^3}{(1-\tilde{\alpha})^3(N-\tilde{\alpha})^3}.
\end{aligned}$$

We want to show that $W(\sigma^{oli,tr}(\tilde{\alpha}))$ is concave in $\tilde{\alpha}$. It is easy to see that

$$\begin{aligned}
\frac{\partial^2 W(\sigma^{oli,tr}(\tilde{\alpha}))}{\partial \tilde{\alpha}^2} &= -(\sigma_2^2+b(\tilde{\alpha})^2)+4(\alpha-\tilde{\alpha})\frac{\partial b(\tilde{\alpha})}{\partial \tilde{\alpha}}b(\tilde{\alpha})+\\
&\quad -[\tilde{\alpha}^2(1-\alpha)+\alpha(1-\tilde{\alpha})^2]\left[\frac{\partial^2 b(\tilde{\alpha})}{\partial \tilde{\alpha}^2}b(\tilde{\alpha})+\left(\frac{\partial b(\tilde{\alpha})}{\partial \tilde{\alpha}}\right)^2\right].
\end{aligned}$$

Using C.10 we can rewrite the above expression as

$$\begin{aligned}\frac{\partial^2 W(\sigma^{oli,tr}(\tilde{\alpha}))}{\partial\tilde{\alpha}^2} &= \frac{4(\alpha-\tilde{\alpha})^2-[\tilde{\alpha}^2(1-\alpha)+\alpha(1-\tilde{\alpha})^2]}{\alpha-\tilde{\alpha}}\frac{\partial b(\tilde{\alpha})}{\partial\tilde{\alpha}}b(\tilde{\alpha})+ \\ &\quad -[\tilde{\alpha}^2(1-\alpha)+\alpha(1-\tilde{\alpha})^2]\left[\frac{\partial^2 b(\tilde{\alpha})}{\partial\tilde{\alpha}^2}b(\tilde{\alpha})+\left(\frac{\partial b(\tilde{\alpha})}{\partial\tilde{\alpha}}\right)^2\right] \\ &= \frac{4(\alpha-\tilde{\alpha})^2-[\tilde{\alpha}^2(1-\alpha)+\alpha(1-\tilde{\alpha})^2]}{\alpha-\tilde{\alpha}}\frac{\partial b(\tilde{\alpha})}{\partial\tilde{\alpha}}b(\tilde{\alpha})+ \\ &\quad -\frac{[\tilde{\alpha}^2(1-\alpha)+\alpha(1-\tilde{\alpha})^2](\alpha-\tilde{\alpha})}{(\alpha-\tilde{\alpha})}\left[\frac{\partial^2 b(\tilde{\alpha})}{\partial\tilde{\alpha}^2}\left(\frac{\partial b(\tilde{\alpha})}{\partial\tilde{\alpha}}\right)^{-1}+\right. \\ &\quad \left.+\left(\frac{\partial b(\tilde{\alpha})}{\partial\tilde{\alpha}}\right)b(\tilde{\alpha})^{-1}\right]\frac{\partial b(\tilde{\alpha})}{\partial\tilde{\alpha}}b(\tilde{\alpha}),\end{aligned}$$

and that we know that at the optimum $\alpha-\tilde{\alpha}>0$ so that the sign of the derivative is the same of the sign of

$$\begin{aligned}&4(\alpha-\tilde{\alpha})^2-\tilde{\alpha}^2(1-\alpha)-\alpha(1-\tilde{\alpha})^2 \\ &\quad -\left[\tilde{\alpha}^2(1-\alpha)+\alpha(1-\tilde{\alpha})^2\right](\alpha-\tilde{\alpha})\left[\frac{\partial^2 b(\tilde{\alpha})}{\partial\tilde{\alpha}^2}\left(\frac{\partial b(\tilde{\alpha})}{\partial\tilde{\alpha}}\right)^{-1}+\left(\frac{\partial b(\tilde{\alpha})}{\partial\tilde{\alpha}}\right)b(\tilde{\alpha})^{-1}\right].\end{aligned}$$

Next, note that $\frac{\partial^2 b(\tilde{\alpha})}{\partial\tilde{\alpha}^2}\left(\frac{\partial b(\tilde{\alpha})}{\partial\tilde{\alpha}}\right)^{-1} > 2\frac{(1+N-2\tilde{\alpha})-\tilde{\alpha}(1-\tilde{\alpha})}{(1-\tilde{\alpha})(N-\tilde{\alpha})} > \frac{2}{1-\tilde{\alpha}}$, and that $\left(\frac{\partial b(\tilde{\alpha})}{\partial\tilde{\alpha}}\right)b(\tilde{\alpha})^{-1} = \frac{N-\tilde{\alpha}^2}{\alpha(N-\tilde{\alpha})(1-\alpha)} > \frac{1}{\tilde{\alpha}(1-\tilde{\alpha})}$. So that it holds that

$$4(\alpha-\tilde{\alpha})^2-\left[\tilde{\alpha}^2(1-\alpha)+\alpha(1-\tilde{\alpha})^2\right]\left[1+(\alpha-\tilde{\alpha})\frac{1+2\tilde{\alpha}}{\tilde{\alpha}(1-\tilde{\alpha})}\right]\leq 0 \implies \frac{\partial^2 W(\sigma^{oli,tr}(\tilde{\alpha}))}{\partial\tilde{\alpha}^2}<0.$$

Standard algebra shows that the expression on the LHS above has the same sign as

$$h(\tilde{\alpha},\alpha)=4\alpha^2\tilde{\alpha}+4\tilde{\alpha}^3-6\alpha\tilde{\alpha}^2-\tilde{\alpha}^4-\alpha^2.$$

when $\tilde{\alpha}$ is at the optimum. Fix $\alpha\in(0,1)$. We start by showing that $h(\tilde{\alpha},\alpha)$ is monotone in $\tilde{\alpha}$. First, note that $h_1(\tilde{\alpha},\alpha)\geq 0 \iff \alpha^2+3\tilde{\alpha}^2-3\alpha\tilde{\alpha}-\tilde{\alpha}^3\geq 0$. Second, note that α^2+

$3\tilde{\alpha}^2 - 3\alpha\tilde{\alpha} - \tilde{\alpha}^3$ is decreasing in $\tilde{\alpha}$ on $[0, 1-\sqrt{1-\alpha})$ and increasing in $\tilde{\alpha}$ on $(1-\sqrt{1-\alpha}, 1]$, while it achieves a minimum at $\tilde{\alpha}_0 = 1-\sqrt{1-\alpha}$. By substitution one can easily verify that $\alpha^2 + 3\tilde{\alpha}_0^2 - 3\alpha\tilde{\alpha}_0 - \tilde{\alpha}_0^3 \geq 0 \iff 2(1-\sqrt{1-\alpha}) - \alpha \geq 0$. But note that $2(1-\sqrt{1-\alpha}) - \alpha \geq 0$ is increasing in α and it is non-negative for both $\alpha = 0$ and $\alpha = 1$, from which follows that $2(1-\sqrt{1-\alpha}) - \alpha \geq 0$, and hence $h_1(\tilde{\alpha}, \alpha) \geq 0$. Given the monotonicity of h_1, it is sufficient to prove that $h(0, \alpha)$ and $h(\alpha, \alpha)$ are both negative. This is easily verified, given that $h(0, \alpha) = -\alpha < 0$ and $h(\alpha, \alpha) = -(1-\alpha)^2 < 0$. This implies that, whenever $\alpha \in (0, 1)$, at the optimum it holds $\frac{\partial^2 W(\sigma^{oli,tr}(\tilde{\alpha}))}{\partial \tilde{\alpha}^2} < 0$.

We use the concavity result to establish the comparative statics of proposition 5 (ii) using the implicit function theorem. First, note that $\frac{\partial^2 W(\sigma^{oli,tr}(\tilde{\alpha}))}{\partial \tilde{\alpha} \partial \sigma_2^2} = \alpha - \tilde{\alpha} > 0$ implies $\frac{\partial \tilde{\alpha}^{oli,tr}}{\partial \sigma_2^2} > 0$. Second, to see that the optimal central banker is increasingly *quailish* the more competitive the banking sector, note that

$$\frac{\partial W(\sigma^{oli,tr}(\tilde{\alpha}))}{\partial \tilde{\alpha} \partial N} = -2(\alpha - \tilde{\alpha}) \frac{(1-\tilde{\alpha})(\tilde{\alpha}\beta)^2}{(1-\tilde{\alpha})^3 (N-\tilde{\alpha})^3} + \left[\tilde{\alpha}^2(1-\alpha) + \alpha(1-\tilde{\alpha})^2\right] \frac{\tilde{\alpha}\beta^2 \left(3\frac{N-\tilde{\alpha}^2}{N-\tilde{\alpha}} - 1\right)}{(1-\tilde{\alpha})^3 (N-\alpha)^3}$$

so that

$$\frac{\partial W(\sigma^{oli,tr}(\tilde{\alpha}))}{\partial \tilde{\alpha} \partial N} > 0 \iff -2(\alpha - \tilde{\alpha})\tilde{\alpha} + \left[\tilde{\alpha}^2 \frac{1-\alpha}{1-\tilde{\alpha}} + \alpha(1-\tilde{\alpha})\right] \left(3\frac{N-\tilde{\alpha}^2}{N-\tilde{\alpha}} - 1\right) > 0$$

But $-2(\alpha - \tilde{\alpha})\tilde{\alpha} + \left[\tilde{\alpha}^2 \frac{1-\alpha}{1-\tilde{\alpha}} + \alpha(1-\tilde{\alpha})\right] \left(3\frac{N-\tilde{\alpha}^2}{N-\tilde{\alpha}} - 1\right) > 2\left[\alpha(1-\tilde{\alpha}) - \tilde{\alpha}(\alpha - \tilde{\alpha})\right] > 0$, since $0 < \tilde{\alpha} < \alpha < 1$. Hence, we have that $\frac{\partial W(\sigma^{oli,tr}(\tilde{\alpha}))}{\partial \tilde{\alpha} \partial N} > 0$, implying $\frac{\partial \tilde{\alpha}^{oli,tr}}{\partial N} > 0$. To see that

$\frac{\partial \tilde{\alpha}^{oli,tr}}{\partial \alpha} > 0$, it is sufficient to prove that $\frac{\partial^2 W(\sigma^{oli,tr}(\tilde{\alpha}))}{\partial \tilde{\alpha} \partial \alpha} < 0$. Note that

$$
\begin{aligned}
\frac{\partial^2 W(\sigma^{oli,tr}(\tilde{\alpha}))}{\partial \tilde{\alpha} \partial \alpha} &= (\sigma_2^2 + b(\tilde{\alpha})^2) - (1-2\tilde{\alpha}) \frac{\partial b(\tilde{\alpha})}{\partial \tilde{\alpha}} b(\tilde{\alpha}) \\
&= \left\{ \left[\frac{\tilde{\alpha}^2(1-\alpha) + \alpha(1-\tilde{\alpha})^2}{\alpha - \tilde{\alpha}} \right] - (1-2\tilde{\alpha}) \right\} \frac{\partial b(\tilde{\alpha})}{\partial \tilde{\alpha}} b(\tilde{\alpha}) \\
&= \left[\frac{\tilde{\alpha}(1-\tilde{\alpha})}{\alpha - \tilde{\alpha}} \right] \frac{\partial b(\tilde{\alpha})}{\partial \tilde{\alpha}} b(\tilde{\alpha}) > 0,
\end{aligned}
$$

so that $\frac{\partial \tilde{\alpha}^{oli,tr}}{\partial \alpha} > 0$ follows from the implicit function theorem.

The main result for $\tilde{\alpha}^{oli}$ is derived analogously, from (C.9). Note that

$$
W(\sigma_{\bar{P}}^{oli}(\tilde{\alpha})) = W(\sigma_{\bar{P}}^{oli,tr}(\tilde{\alpha})) - \frac{1}{2}[\tilde{\alpha}^2(1-\alpha) + \alpha(1-\tilde{\alpha})^2] \hat{\sigma}_{1,\bar{P}}^2(\tilde{\alpha}).
$$

Existence of $al\tilde{p}ha^{oli}$ is guaranteed by the continuity of $\hat{\sigma}_{1,\bar{P}}^2(\tilde{\alpha})$. For $al\tilde{p}ha^{oli} \in (0, \alpha)$ to hold, it is sufficient to show that

$$
-\frac{1}{2}[\tilde{\alpha}^2(1-\alpha) + \alpha(1-\tilde{\alpha})^2] \hat{\sigma}_{1,\bar{P}}^2(\tilde{\alpha})
$$

is weakly decreasing in $\tilde{\alpha}$ if $\tilde{\alpha} \geq \alpha$ and weakly increasing in $\tilde{\alpha}$ in a neighborhood of $\tilde{\alpha} = 0$. First, given that $\hat{\sigma}_{1\bar{P}}(\tilde{\alpha})$ is always weakly increasing in $\tilde{\alpha}$ and $\tilde{\alpha}^2(1-\alpha) + \alpha(1-\tilde{\alpha})^2$ is weakly increasing in $\tilde{\alpha}$ if $\tilde{\alpha} \geq \alpha$, we have that $\tilde{\alpha} < \alpha$. Second, the right derivative of $\hat{\sigma}_{1\bar{P}}(\tilde{\alpha})$ at $\alpha = 0$ is equal to 0, which guarantees $\tilde{\alpha}^{oli} > 0$.

Finally, we prove part (iii). First, to see that under transparent communication markets are worse off than under transparency, note that market players' payoff as a function of $\tilde{\alpha}$ are

$$
\begin{aligned}
EU_i^{oli,tr} &= -\frac{1}{2}(1-\tilde{\alpha})^2(\sigma_2^2 + \beta(\tilde{\alpha})^2) + \frac{(\tilde{\alpha}\beta)^2}{(N-\tilde{\alpha})(1-\tilde{\alpha})} \\
&= -\frac{1}{2}(1-\tilde{\alpha})^2\sigma_2^2 + \frac{(2N-1-\tilde{\alpha})(\tilde{\alpha}\beta)^2}{2(N-\tilde{\alpha})^2(1-\tilde{\alpha})} \\
&= -\frac{1}{2}(1-\tilde{\alpha})^2\sigma_2^2 + \left(\frac{\tilde{\alpha}\beta}{N-\tilde{\alpha}}\right)^2 \left(\frac{N-1}{1-\tilde{\alpha}} + \frac{1}{2}\right)
\end{aligned}
$$

which is clearly increasing in $\tilde{\alpha}$. To see that under cheap talk communication markets can be better off with a kitish central banker relative to an unbiased one, it is sufficient to find an example. The example is provided at the end of section 3.4.2.

Proof of proposition 6

We want to show that there exist a $\delta^\star \in (0,1)$ such that the proposed strategy is a PBE of the repeated game with discount rate $\delta \geq \delta^\star$. Note that in the punishment phase players revert to a fixed stage game PBE, so that it is sufficient to check that, on the path of play, no player has incentive to deviate.

Consider *CB* first. Note that *CB* actions at time $s < \tau$ do not influence period τ play of any player, so it is sufficient to check that *CB* has no way to increase its stage payoff via a deviation. The latter statement follows immediately from the fact that stage game play on the path of play satisfy the efficiency conditions of 2.

Turning to the deviation incentives of player $i \in I$, the no profitable deviation condition requires

$$\frac{1}{2}\left[\frac{\alpha\beta}{N-\alpha}\right]^2 + \frac{\delta}{1-\delta}\left\{\frac{1}{2}\left[\frac{\alpha\beta}{N-\alpha}\right]^2\frac{2N-1-\alpha}{(1-\alpha)} - \frac{1}{2}(1-\alpha)^2 Var(\omega_1)\right\} \leq 0,$$

where the left hand side is the net benefit that i obtains if she deviates to the stage best response once and then follows the equilibrium strategy in all future stages. Using $Var(\omega_1) = \frac{\phi_1^2}{3}$, it is easily seen from the above condition that

$$\delta^\star \equiv \frac{\left[\frac{\alpha\beta}{N-\alpha}\right]^2}{\left[\frac{\alpha\beta}{N-\alpha}\right]^2\left[1 - \frac{2N-1-\alpha}{1-\alpha}\right] + \frac{\phi_1^2}{3}(1-\alpha)^2}.$$

For the equilibrium to exist it must be that $\delta^\star < 1$, which is true if and only if
$\phi_1 > \frac{\alpha\beta}{(N-\alpha)(1-\alpha)}\sqrt{\frac{2N-1-\alpha}{1-\alpha}}$.

Proof of proposition 7

For part (i), it is sufficient to provide a collusive *PBE* strategy profile that implements efficient play on path in every stage, that is, such that stage play on the equilibrium path satisfies the conditions of 2. We will then verify that such equilibrium exists if $\hat{\sigma}^2_{1,\bar{P}}(N) > \left[\frac{\alpha\beta}{(N-\alpha)(1-\alpha)}\right]^2 \frac{2N-\alpha-1}{1-\alpha}$.

Consider the following profile. At $\tau = 0$, *CB* chooses $m^\tau = \omega_1^\tau$; at $\tau > 0$, *CB* select $m^\tau = \omega_1^\tau$ if $x_i^s = m^s$ for each $s < \tau$ and $i \in I$; if there exists $s < \tau$ and $i \in I$ such that $x_i^s \neq m^s$, then *CB* uses the a stage communication rule $\sigma_{m,\bar{P}}$ corresponding to the most informative stage game PBE when the bias is at the no-coordination $\frac{\alpha\beta}{(N-\alpha)(1-\alpha)}$. Each investor $i \in I$ plays the following strategy. At $\tau = 0$, $x_i^\tau = m^\tau$. At $\tau > 0$, $i \in I$ selects $x_i^\tau = m^\tau$ if $x_j^s = m^s$ for each $s < \tau$ and $j \in I$; if there exists $s < \tau$ and $j \in I$ such that $x_j^s \neq m^s$, then $x_i^\tau = \mathbb{E}[\omega_1^\tau | m^\tau, \sigma_{m,\bar{P}}] - \frac{\alpha\beta}{(N-\alpha)(1-\alpha)}$. At each $\tau \geq 0$, *CB* uses the policy rule $r^\tau = (1-\alpha)\omega^\tau + \alpha\bar{x}^\tau$.

By construction, in every stage the *CB* communicates as much as possible given the equilibrium expected investment bias of that period. As in the case of the previous proposition, punishment is carried out by playing a fixed stage Nash in any period after a deviation (regardless of how the history of play evolves after the deviation). Hence deviations are not profitable during the punishment phase. Moreover, the *CB* has no incentive to deviate from the equilibrium path, since efficient play is implemented on the equilibrium path in each stage. Hence, for the strategy profile considered to be a PBE of the repeated game it is sufficient to impose that $i \in I$ has no incentive to make a one shot deviation from the equilibrium path. This requires

$$\frac{1}{2}\left[\frac{\alpha\beta}{N-\alpha}\right]^2 + \frac{\delta}{1-\delta}\left\{\frac{1}{2}\left[\frac{\alpha\beta}{N-\alpha}\right]^2 \frac{2N-1-\alpha}{1-\alpha} - \frac{1}{2}(1-\alpha)^2\hat{\sigma}^2_{1,\bar{P}}(N)\right\} \leq 0,$$

where the left hand side is the net benefit that i obtains if she does a one-shot deviation towards the stage best response and then follows the equilibrium strategy in all future stages. Note that at

$\delta = \delta_1^\star$ the previous equation must hold with equality. Simple algebra yields

$$\delta_1^\star \equiv \frac{\left[\frac{\alpha\beta}{N-\alpha}\right]^2}{\left[\frac{\alpha\beta}{N-\alpha}\right]^2\left[1-\frac{2N-1-\alpha}{1-\alpha}\right]+(1-\alpha)^2\hat{\sigma}^2_{1,\bar{P}}(N)}.$$

For the equilibrium to exist it must be that $\delta^\star < 1$, which is true if and only if $\hat{\sigma}^2_{1,\bar{P}}(N) > \left[\frac{\alpha\beta}{(N-\alpha)(1-\alpha)}\right]^2 \frac{2N-1-\alpha}{1-\alpha}$. Note that the latter inequality also guarantees that i's stage-game payoff under the equilibrium considered is greater than under the most informative stage-game equilibrium for the parameters considered.

For part (ii), it is sufficient to provide a collusive *PBE* strategy profile that mimics the stage game equilibrium when $N = 1$ on the path of play in every stage. We will then verify that the existence of such equilibrium requires $\hat{\sigma}^2_{1,\bar{P}}(1) - \hat{\sigma}^2_{1,\bar{P}}(N) < \left[\frac{\alpha\beta}{(1-\alpha)^2}\right]^2$.

Consider the following profile. At $\tau = 0$, *CB* uses a stage communication rule $\sigma^1_{m,\bar{P}}$ corresponding to the most informative stage game PBE when investors' bias is at the monopolistic level $\frac{\alpha\beta}{(1-\alpha)^2}$; at $\tau > 0$, *CB* keeps using the same communication rule if $x_i^s = \mathbb{E}[\omega_1^s|m^s,\sigma^1_{m,\bar{P}}] - \frac{\alpha\beta}{(1-\alpha)^2}$ for each $s < \tau$ and $i \in I$; if there exists $s < \tau$ and $i \in I$ such that $x_i^s \neq m^s$, then *CB* uses a stage communication rule $\sigma^N_{m,\bar{P}}$ corresponding to the most informative stage game PBE when the bias is at the no-coordination level $\frac{\alpha\beta}{(N-\alpha)(1-\alpha)}$. Each investor $i \in I$ plays the following strategy. At $\tau = 0$, $x_i^\tau = \mathbb{E}[\omega_1^\tau|m^\tau,\sigma^1_{m,\bar{P}}] - \frac{\alpha\beta}{(1-\alpha)^2}$. At $\tau > 0$, $i \in I$ selects $x_i^\tau = \mathbb{E}[\omega_1^\tau|m^\tau,\sigma^1_{m,\bar{P}}] - \frac{\alpha\beta}{(1-\alpha)^2}$ if $x_j^s = m^s$ for each $s < \tau$ and $j \in I$; if there exists $s < \tau$ and $j \in I$ such that $x_j^s \neq \mathbb{E}[\omega_1^s|m^s,\sigma^1_{m,\bar{P}}] - \frac{\alpha\beta}{(1-\alpha)^2}$, then $x_i^\tau = \mathbb{E}[\omega_1^\tau|m^\tau,\sigma^N_{m,\bar{P}}] - \frac{\alpha\beta}{(N-\alpha)(1-\alpha)}$. At each $\tau \geq 0$, *CB* uses the policy rule $r^\tau = (1-\alpha)\omega^\tau + \alpha\bar{x}^\tau$.

By construction, in every stage the *CB* communicates as much as possible given the equilibrium expected investment bias of that period. As in the previous cases, punishment is carried out by playing a fixed stage Nash in any period after a deviation (regardless of how the history of play evolves after the deviation). Hence deviations are not profitable during the punishment phase. The *CB* has no incentive to deviate from the equilibrium path, since it is maximizing its stage game expected payoff given the investors' strategy, and *CB* deviations have no influence on subsequent play. Hence, for the strategy profile outlined to be a PBE of the repeated game it is sufficient to impose that $i \in I$ has no incentive to make a one-shot deviation from the equilibrium path. Following the same procedure as in the previous proofs one can easily show that $\delta_2^\star$ exists in $(0,1)$ as long as the punishment stage payoff of investor $\imath \in I$ is strictly lower than her on path stage payoff. The condition requires

$$-\frac{1}{2}(1-\alpha)^2(\hat{\sigma}^2_{1,\bar{P}}(1)+\sigma_2^2)+\frac{1}{2}\left[\frac{\alpha\beta}{1-\alpha}\right]^2 > -\frac{1}{2}(1-\alpha)^2(\hat{\sigma}^2_{1,\bar{P}}(N)+\sigma_2^2)+\frac{1}{2}\left[\frac{\alpha\beta}{N-\alpha}\right]^2\frac{2N-1-\alpha}{1-\alpha},$$

or equivalently,

$$\hat{\sigma}^2_{1,\bar{P}}(1)-\hat{\sigma}^2_{1,\bar{P}}(N) < \left[\frac{\alpha\beta}{(1-\alpha)^2}\right]^2-\left[\frac{\alpha\beta}{(N-\alpha)(1-\alpha)}\right]^2\frac{2N-1-\alpha}{1-\alpha}.$$

which simply amounts to requiring that the policy influence gain thanks to the exercise of monopolistic market power is above the surprise loss due to less transparent communication than in the no-coordination case. Letting $N \to \infty$ in the previous expression, it is verified that markets are better off in the monopolistic equilibrium relative to the efficient equilibrium if and only if $\hat{\sigma}^2_{1,\bar{P}}(1) < \left[\frac{\alpha\beta}{(1-\alpha)^2}\right]^2$, which completes the proof.

C.2 Silicon Valley Bank and Signature Bank Case Study

In this Appendix, we provide a case study of Silicon Valley Bank (SVB) and Signature Bank (SB) to help illuminate a key idea in our theory: when the central bank is concerned about financial sector losses, it will underreact to inflationary pressure or output gap concerns (see Proposition 1). This underreaction will take the form of interest rates being lower than they would have been had the central bank not been concerned by such losses. In the language of our model, the interest rate will be lower than the economy-stabilizing rate (i.e., the rate needed for inflation/output gap stability).[4]

SVB was a bank with over $200 billion in assets and SB was a bank with over $100 billion in assets. In March 2023, SVB faced severe problems as a result of interest rate risk, which it had actively not hedged (Metrick (2024)). The increases in interest rates combined with inadequate hedging negatively impacted its asset valuations which led to its failure when depositors began to withdraw their deposits. The failure of SVB led to widespread concerns and similarly rapid deposit outflows at SB which subsequently failed. These failures ultimately led to the '2023 Banking Crisis' and broader concerns around systemic risk.[5]

The pertinent question from the perspective of our theoretical model is whether these concerns around systemically risky banking sector losses induced the Fed to underreact to the high inflation during that period. Specifically, our theory predicts that the Fed set the fed funds rate (FFR) lower than it would have done in the counterfactual world of no banking crisis. However, this is difficult to prove given the empirical challenge of knowing that counterfactual.

We use three different methods to consider the counterfactual world: (1) Taylor rule predictions; (2) Fed Funds futures implied probabilities; and (3) FOMC minutes. Each of these support the idea that the Fed kept rates lower than it would have done in the absence of banking

[4]Our theory also predicts that knowing the central bank will underreact, financial institutions will take greater risk. This case study is only intended to show an example of the central bank's underreaction rather than the consequent risk-taking decisions of financial institutions.

[5]https://www.fdic.gov/news/press-releases/2023/pr23017.html

sector concerns and therefore are consistent with the predictions from our theoretical model. Below, we describe each method in more detail.

Taylor rule predictions

The Taylor rule is an equation that prescribes a value for the FFR based on the values of inflation relative to target and the output gap (Taylor (1993)). In a sense, it captures the typical central bank loss function (the first term in equation (3.2)) and can be thought of as a theory-based prediction. The Taylor rule (and modifications of it) have been used consistently as a 'benchmark' for the FFR. Therefore, one can compare the actual FFR to that predicted by the Taylor rule to see whether the Fed is keeping rates too low or too high.

We estimate the following Taylor rule[6]:

$$\hat{r}_t = \bar{r}_t + \pi_t + \theta(\pi_t - \pi_t^*) + (1-\theta)\tilde{Y}_t \tag{C.11}$$

where $\hat{r}_t$ is the predicted FFR, $\bar{r}_t$ is the long-run real interest rate, π_t is inflation, π_t^* is the target inflation rate, $\tilde{Y}_t$ is the output gap, and θ is the weight the central bank puts on inflation. We use commonly used values for the different variables. Specifically, we set the long-run real interest rate to 2%, the target inflation rate to 2%, and θ to 0.5.

Given the underlying data is quarterly, we can look at the predicted rate versus the actual rate in the second quarter of 2023 (the first data point after the Fed's rate-setting meeting). Our theory suggests that the Taylor rule predicted rate should be higher than the Fed's chosen rate because the former does not account for banking sector losses. Indeed, (C.11) predicts a FFR of 6.17% while the actual FFR was only 4.83%. While this is consistent with our theory, the Taylor rule is typically not an accurate predictor of the FFR and therefore may not serve as effective counterfactual.

[6]Note that we use a simple Taylor rule. One can use a number of different modified rules. For the purposes of our analysis, the different rules overall do not make a material difference.

Fed Funds futures implied probabilities

Next, we consider a market-based prediction of the FFR. Specifically, the probabilities of different rate decisions are computed using 30-day Fed Funds futures pricing data.[7] The useful feature of this method is that we can see the distribution of expectations rather than a single outcome.

Figure C.1 below shows the probabilities of the different actions that the Fed will take according to market expectations. We look at the different probabilities from Friday 10 March to Friday 17 March. The important events are that between Friday 10 March and Sunday 12 March, both SVB and SB failed. Therefore, our theory would predict that FFR hikes should be less likely after that weekend than before it. This is precisely what we see. On Friday 10 March, the market priced in a 60% chance of a 25 basis point hike, a 40% chance of a 50 basis point hike, and a 0% chance of no change. These were driven by the very high inflation at the time. However, just a few days later, there was a significant change in these probabilities. The probability of a 50 basis point hike fell to zero percent and stayed there while the probability of no change rose to over a third.

[7] Our data is from the CME FedWatch Tool.

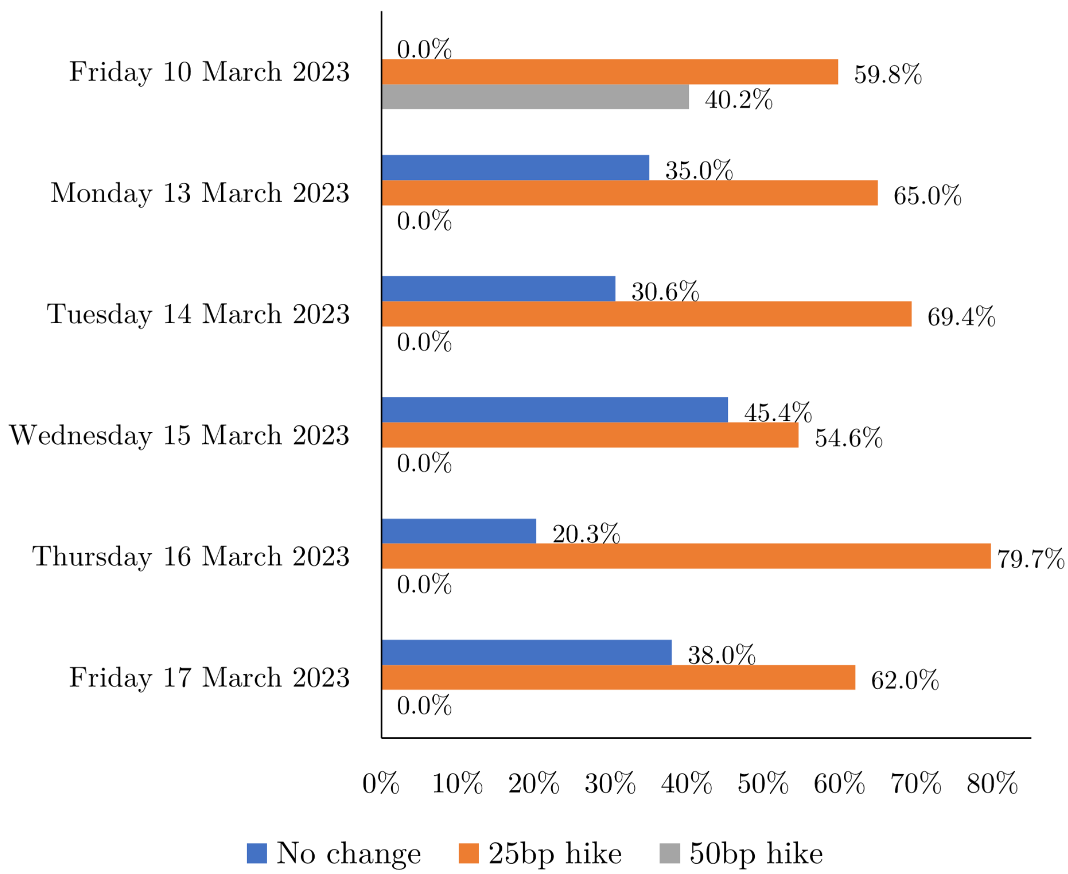

Figure C.1. Probability of different FFR actions

The above Figure is consistent with our theory as inflation concerns did not suddenly disappear over that weekend. However, it is still possible, albeit unlikely, that the Fed may have been concerned about something else and that this timing is merely a coincidence.

FOMC Minutes

Finally, to show that the Fed was explicitly concerned about banking sector developments and this played a role in its decision to not raise rates, we examine the minutes of the FOMC meeting in March 2023. The following excerpt from the minutes provide much clearer evidence of a counterfactual:

"Some participants noted that given persistently high inflation and the strength of recent economic data, they would have considered a 50 basis point increase in the target range to have been appropriate at this meeting *in the absence of recent developments in the banking sector* [emphasis added]."

The Fed ultimately decided on a 25 basis point hike which is consistent with our theory that the Fed did not raise rates as much as it would have done had there been no banking sector concerns.

Ultimately, we believe that these three pieces of evidence show, through a leading example, how systemic stresses in the banking sector can lead that Fed to underreact to inflation or the output gap, consistent with our theoretical model.[8]

[8]Note that to have rigorous empirical evidence, one would need to conduct a systematic study of the data rather than a single case study.

Appendix D

Supplemental Material to Chapter 4

D.1 Proof of Proposition 1

We begin by defining the domar weights of the economy in terms of primitives. Using the labour market clearing condition to obtain:

$$\frac{p_i y_i}{Y} = \beta_i + \mu \sum_j a_{ji} \frac{p_i y_i}{Y}.$$

Denoting $\frac{p_f y_f}{Y} = \lambda_f$, this implies the following Domar weights:

$$(I - \mu A)\lambda = \beta \Leftrightarrow \lambda = (I - \mu A)^{-1}\beta$$

Next, we totally differentiate the balance sheet of banks which yields:

$$\begin{aligned} -\theta d\ln R_b + \textstyle\sum_i \frac{s_{ib}}{\sum_i s_{ib}} \left[d\ln w + d\ln l_f + \theta \textstyle\sum_{b'} s_{ib'} d\ln R_{b'} \right] \\ = I_{(1-\alpha_b^e)} \left[d\ln w + \rho d\ln R_{bd} - \textstyle\sum_{b'} \frac{\phi_{b'} R_{b'd}^{\rho}}{\sum_k \phi_k R_{kd}^{\rho}} d\ln R_{bd} \right] + I_{\alpha_b^e} d\ln e_b. \end{aligned}$$

Now, we use the fact that $wl_i R_i = (1-\mu)\lambda_i Y$ to obtain:

$$d\ln w + d\ln l_i = d\ln Y - d\ln R_i.$$

The change in each firm's financing rate, in turn, equals,

$$d\ln R_f = \sum_b s_{fb} d\ln R_b$$

Plugging into the total derivative of the balance sheet equation, we obtain

$$\begin{aligned} & -\theta d\log R_b + \sum_i \omega_{bi} \left(d\ln Y + (\theta-1) \sum_{b'} s_{ib'} \hat{R}_{b'} \right) \\ & = \alpha_b^e + \alpha_b^{\mathscr{D}} \left[d\ln w + \rho d\ln R_{bd} - \sum_{b'} \frac{\phi_{b'} R_{b'd}^{\rho}}{\sum_k \phi_k R_{kd}^{\rho}} d\ln R_{bd} \right] \end{aligned}$$

where $\omega_{bi} = \frac{s_{ib}}{\sum_i s_{ib}}$

Simplifying:

$$\begin{aligned} & -\theta d\log R_b + (\theta-1) \sum_i \omega_{bi} \sum_{b'} s_{ib'} \hat{R}_{b'} \\ & = \alpha_b^{\mathscr{D}} d\ln w - d\ln Y + \alpha_b^{\mathscr{D}} \left[\rho d\ln R_{bd} - \sum_{b'} \gamma_{b'} d\ln R_{b'd} \right] + \alpha_b^e d\ln e_{bd} \end{aligned}$$

Stacking $\hat{R} = \{d\ln R_b\}_b$ and $\hat{\mathbf{R}}_{\mathscr{D}} = \{d\ln R_{bd}\}_b$, this can be written

$$\begin{aligned} & [-\theta I_{B\times B} + (\theta-1) \times \Omega \times S]\, \hat{\mathbf{R}} \\ & = I_{1-\alpha^e} \hat{w} - I_{B\times 1} \hat{Y} + \rho \times I_{\alpha_b^{\mathscr{D}}} \times \hat{\mathbf{R}}_{\mathscr{D}} - 1_{B\times 1} \otimes \gamma' \odot \left(\alpha'_{\mathscr{D}} + 1_{1\times B} \right) \hat{\mathbf{R}}_{\mathscr{D}} + I_{\alpha^e} \hat{\mathbf{e}}, \end{aligned}$$

where $I_{\alpha^e} = Diag(\{\alpha_b^e\}_b)$, $\Omega \equiv \{\omega_{bi}\}_{b\in\mathscr{B}, i\in\mathscr{I}}$ is the $B \times N$ matrix of bank-firm relationships and $S \equiv \{s_{ib}\}_{i\in\mathscr{I}, b\in\mathscr{B}}$ is the $N \times B$ matrix of firm-bank relationships in the loan market. The assumption of complete pass-through of loan into deposit rates then implies that

$$\hat{R} = \Gamma\left[I_{1-\alpha^e}\hat{w} - I_{1-\alpha^e}\hat{Y} + I_{\alpha^e}\hat{\mathbf{e}}\right]$$

where Γ is the PE pass-through of equity shocks into interest rates, holding fixed GDP and wages.

$$\Gamma \equiv -\left[\theta I_{B\times B} + \rho I_{\alpha_b^{\mathscr{D}}} - (\theta-1)\times\Omega\times S - 1_{B\times 1}\otimes\gamma'\odot\left(\alpha'_{\mathscr{D}} + 1_{1\times B}\right)\right]^{-1}$$

To solve for the change in the wage, we write the change in each firm's profit-maximising price as follows:

$$d\log\frac{p_i}{w} = \mu d\ln R_i + \mu\sum_k a_{kf} d\log\frac{p_k}{w}.$$

Recalling that $\hat{\mathbf{R}} = S\times\hat{\mathbf{R}}$, we can rewrite this as:

$$\hat{p}_f/\hat{w} = \mu L S\hat{\mathbf{R}}$$

where $L = [1-\mu A']^{-1}$ is the Leontief inverse. Multiplying by $\beta' = \{\beta_f\}'_f$, imposing the normalisation $\beta_f d\ln p_f = 0$, and using the definition of the Domar weights, $\lambda' = \beta L$, we then obtain:

$$\hat{w} = -\mu\lambda' S\hat{R},$$

Next, we calculate profits of banks:

$$\pi_b = L_b R_b - \mathscr{D}_b R_{db} = \frac{\frac{\theta}{\theta-1}}{\frac{\rho}{\rho+1}}(\mathscr{D}_b + e_b)R_{db} - \mathscr{D}_b R_{db} = \left(\left[\frac{\frac{\theta}{\theta-1}}{\frac{\rho}{\rho+1}} - 1\right]\mathscr{D}_b + \frac{\frac{\theta}{\theta-1}}{\frac{\rho}{\rho+1}} e_b\right)R_{db}$$

Aggregating over all banks, we then find:

$$\sum \pi_b = w\left[\frac{\frac{\theta}{\theta-1}}{\frac{\rho}{\rho+1}} - 1\right]\sum \gamma_b R_{db} + \frac{\frac{\theta}{\theta-1}}{\frac{\rho}{\rho+1}}\sum e_b R_{db}$$

GDP can then be written as:

$$\begin{aligned}
Y &= wR_{\mathscr{D}} + \sum_b \pi_b \\
&= w\sum_b \gamma_b R_{db} + \left[\frac{\frac{\theta}{\theta-1}}{\frac{\rho}{\rho+1}} - 1\right]\sum w\gamma_b R_{db} + \frac{\frac{\theta}{\theta-1}}{\frac{\rho}{\rho+1}}\sum e_b R_{db} \\
&= \frac{\frac{\theta}{\theta-1}}{\frac{\rho}{\rho+1}} wR_{\mathscr{D}} + \frac{\frac{\theta}{\theta-1}}{\frac{\rho}{\rho+1}}\sum_b e_b R_{db}
\end{aligned}$$

Define $\omega^Y_w = \frac{\frac{\theta}{\theta-1}}{\frac{\rho}{\rho+1}}\frac{wR_{\mathscr{D}}}{Y}$, and $\iota = [\frac{e_b R_{db}}{\sum_b e_b R_{db}}]_b$. Totally differentiating the above expression for GDP then yields:

$$\hat{Y} = \omega^Y_w\left[\hat{w} + \gamma' \hat{R}\right] + \left(1 - \omega^Y_w\right)\iota'\left[\hat{e} + \hat{R}\right].$$

Combining with the change in the real interest rate and substituting the change in the wage yields

$$\begin{aligned}
&\hat{\mathbf{R}} = \Gamma\left[I_{\alpha^{\mathscr{D}}_b}\hat{w} - I_{B\times 1}\hat{Y} + I_{\alpha^e}\hat{\mathbf{e}}\right] \\
&\Leftrightarrow \left[\Gamma^{-1} + (I_{1-\alpha^e} + \omega^Y_w)\mu\lambda' S + I_{B\times 1}((1-\omega^Y_\Pi)\gamma' + \omega^Y_\Pi \iota')\right]\hat{R} + \left[\Gamma I_{\alpha^e} - \omega^Y_\Pi \iota'\right]\hat{\mathbf{e}} \\
&\Leftrightarrow \hat{R} = \left[\Gamma^{-1} + \left((I_{1-\alpha^e} + \omega^Y_w)\mu\lambda' S + I_{B\times 1}((1-\omega^Y_\Pi)\gamma' + \omega^Y_\Pi \iota')\right)\right]^{-1}\left[\Gamma I_{\alpha^e} - \omega^Y_\Pi \iota'\right]\hat{\mathbf{e}} \\
&\Leftrightarrow \hat{R} \equiv \mathscr{M}\left(\Gamma I_{\alpha^e} - \omega^Y_\Pi \iota'\right)\hat{\mathbf{e}}
\end{aligned}$$

where $1 - \omega^Y_w = \omega^Y_\Pi$ and denotes the share of bank profits in GDP. Substituting back into the change in GDP, we then obtain

$$\hat{Y} = \left(1 - \omega_{\Pi}^{Y}\right) \hat{w} + \omega_{\Pi}^{Y} \iota' \hat{e} + \left[\left(1 - \omega_{\Pi}^{Y}\right) \gamma' + \omega_{\Pi}^{Y} \iota'\right] \hat{R}$$

$$\Leftrightarrow \hat{Y} = -\left[\omega_{w}^{Y} \mu \lambda' S - \omega_{w}^{Y} \gamma' - \left(1 - \omega_{w}^{Y}\right) \iota'\right] \mathscr{M} \left[\Gamma \alpha_{b}^{e} - (1 - \omega_{w}^{Y}) I_{B\times 1} \iota'\right] \hat{\mathbf{e}} + \omega_{\Pi}^{Y} \iota' \hat{e}$$

$$\therefore d \log Y = -\left[(1 - \omega_{\Pi}^{Y})(\mu \lambda' S - \gamma') - \omega_{\Pi}^{Y} \iota'\right] \mathscr{M} \left[\Gamma \alpha_{b}^{e} - \omega_{\Pi}^{Y} I_{B\times 1} \iota'\right] \hat{\mathbf{e}} + \omega_{\Pi}^{Y} \iota' \hat{\mathbf{e}}$$

This completes the proof.