

UNIVERSITY OF BERN

ESSAYS IN MACROECONOMICS

Lucas Christopher Andreas Kyriacou

Inauguraldissertation zur Erlangung der Würde eines

DOCTOR RERUM OECONOMICARUM

der Wirtschafts- und Sozialwissenschaftlichen Fakultät

Januar 2023

Originaldokument gespeichert auf dem Webserver der Universitätsbibliothek Bern.



Dieses Werk ist unter einem Creative Commons
Namensnennung - Nicht-kommerziell - Keine Bearbeitung 2.5 Schweiz
(CC BY-NC-ND 2.5 CH) Lizenzvertrag lizenziert. Um die Lizenz anzusehen, gehen Sie bitte
auf <http://creativecommons.org/licenses/by-nc-nd/2.5/ch/>.

Urheberrechtlicher Hinweis

Dieses Werk steht unter einer Lizenz der Creative Commons
Namensnennung - Nicht-kommerziell - Keine Bearbeitung 2.5 Schweiz
(CC BY-NC-ND 2.5 CH)

<http://creativecommons.org/licenses/by-nc-nd/2.5/ch/>.

Sie dürfen:



Teilen - das Material in jedwedem Format oder Medium vervielfältigen und weiterverbreiten.

Unter folgenden Bedingungen:



Namensnennung - Sie müssen angemessene Urheber- und Rechteangaben machen, einen Link zur Lizenz beifügen und angeben, ob Änderungen vorgenommen wurden. Diese Angaben dürfen in jeder angemessenen Art und Weise gemacht werden, allerdings nicht so, dass der Eindruck entsteht, der Lizenzgeber unterstütze gerade Sie oder Ihre Nutzung besonders.



Nicht kommerziell - Sie dürfen das Material nicht für kommerzielle Zwecke nutzen.



Keine Bearbeitungen - Wenn Sie das Material remixen, verändern oder darauf anderweitig direkt aufbauen, dürfen Sie die bearbeitete Fassung des Materials nicht verbreiten.

Im Falle einer Verbreitung müssen Sie anderen die Lizenzbedingungen, unter welche dieses Werk fällt, mitteilen.

Jede der vorgenannten Bedingungen kann aufgehoben werden, sofern Sie die Einwilligung des Rechteinhabers dazu erhalten.

Diese Lizenz lässt die Urheberpersönlichkeitsrechte nach Schweizer Recht unberührt.

Eine ausführliche Fassung des Lizenzvertrags befindet sich unter
<http://creativecommons.org/licenses/by-nc-nd/2.5/ch/legalcode.de>

The faculty accepted this thesis on the 30. March 2023 at the request of the reviewers Prof. Dr. Luca Benati, Prof. Dr. Harris Dellas and Prof. Dr. Fabrice Collard as dissertation, without wishing to comment on the views expressed therein.

Die Fakultät hat diese Arbeit am 30. März 2023 auf Antrag der Gutachter Prof. Dr. Luca Benati, Prof. Dr. Harris Dellas und Prof. Dr. Fabrice Collard als Dissertation angenommen, ohne damit zu den darin ausgesprochenen Auffassungen Stellung nehmen zu wollen.

Contents

Contents	viii
List of Figures	x
List of Tables	xii
Introduction	xiii
1 Identification of Confidence Shocks	1
1.1 Introduction	2
1.2 Assessing Existing Approaches to Identifying Confidence Shocks	5
1.3 Do Survey Measures of Confidence Granger-Cause Macro Variables?	13
1.4 Methodology	15
1.5 Evidence	27
1.6 Conclusion	32
1.A Appendix I	33
2 Business-Cycle Anatomy around the World	37
2.1 Introduction	38
2.2 Literature	39
2.3 Empirical Strategy	41
2.4 Cross-Country Analysis	47
2.5 Opening Up	65
2.6 Conclusion	69
2.A Appendix I	71
2.B Appendix II	72
2.C Appendix III	79
3 Costs of Leaning against the Wind: Evidence from High-Frequency Identification	93

3.1	Introduction	94
3.2	Literature	95
3.3	Data	98
3.4	Empirical Strategy	99
3.5	Evidence	105
3.6	Conclusion	110
3.A	Appendix I	111
3.B	Appendix II	112
4	A Model United Nations Experiment on Climate Negotiations	115
4.1	Introduction	116
4.2	The MUN Experiment: Design	120
4.3	Hypotheses	124
4.4	Experimental Procedures	125
4.5	Results	127
4.6	Conclusion	134
4.A	Appendix I	136
	Bibliography	137
	Statement of Authorship	145

List of Figures

1.1	Spuriously identifying confidence shocks in SW's model via Beaudry et al.'s (2011) identification strategy: fractions of FEV explained by the identified confidence shock	8
1.2	Spuriously identifying confidence shocks in the standard New Keynesian model via sign restrictions: fractions of FEV explained by the identified confidence shock	12
1.3	Experiment 1, IRF	21
1.4	Experiment 1, FEV	22
1.5	Experiment 2, IRF	23
1.6	Experiment 2, FEV	24
1.7	Experiment 3, IRF	25
1.8	Experiment 3, FEV	26
1.9	Real data: Conference Board Overall Index	29
1.10	Real data: Michigan Index 10	30
1.11	Estimated fractions of FEV jointly explained by permanent IS and N shocks, and residual fraction of FEV	31
2.1	The various facets of the MBC shock, IRFs (VECM1)	49
2.2	IRFs to the MBC shocks around the world (VECM1)	52
2.3	The MBC shock and capacity utilization, IRFs (VECM1)	55
2.4	The various facets of the MLR shock, IRFs (VECM1)	60
2.5	Cross-country correlation of output: unconditional and conditional (VECM1) . .	62
2.6	Correlation of actual and counterfactual output with the corresponding US series (VECM1)	63
2.7	IRFs to the MBC shocks around the world (VECM1)	65
2.8	Comparison of MBC shock with import, export, and terms of trade shocks, IRFs (VECM1)	68
2.9	Comovement in key macroeconomic series around the world	74
2.10	Comparison of hours	77
2.11	Comparison of datasets for the US.	78

2.12	The various facets of the MBC shock, IRFs (VECM1)	81
2.13	The various facets of the MBC shock, IRFs (VAR)	83
2.14	IRFs to the MBC shocks around the world (VAR)	84
2.15	The various facets of the MLR shock, IRFs (VAR)	85
2.16	IRFs to the MBC shocks around the world (VECM1)	89
2.17	IRFs to the MBC shocks around the world (VAR)	89
2.18	Comparison of MBC shock with import, export and terms of trade shocks, IRFs (VAR)	91
3.1	IRFs with raw versus orthogonalized surprises	104
3.2	IRFs to a monetary policy shock in the housing dataset	106
3.3	IRFs to a monetary policy shock in the credit dataset	107
3.4	IRFs with raw versus orthogonalized surprises	112
3.5	IRFs to a monetary policy shock in the credit dataset	113
4.1	Actual worldwide reductions of emissions (weighted average over all countries) in Committee C1 versus Committee C2 averaged over all COP meetings.	127
4.2	Resolutions vs. actual reductions. The left bars depict the reductions of emis- sions as announced in the resolutions, the right bars the actual reductions of emissions in all countries (national commitments weighted by country size).	129
4.3	Proclaimed goal vs. actual reductions of each country in C1. The left bars depict the average announced reduction in those resolutions that the country voted for, the right bars depict the average actual reductions in all cases in which the country voted for the resolution.	130
4.4	Fraction of countries that voted in favor of the resolution in Committee C1 and Committee C2 (averaged over all COP meetings).	131
4.5	Reduction of emissions of each country in Committee C1 and Committee C2 in 2030-2050 (averaged over all COP meetings).	133

List of Tables

1.1	Zero and sign restrictions from Beaudry, Nam, et al. (2011) and Nam and Wang (2019)	5
1.2	Wald tests for Granger-causality of consumer confidence indices from the Conference Board onto macroeconomic variables, and bootstrapped p-values ^a	14
1.3	Wald tests for Granger-causality of consumer confidence indices from the Michigan survey onto macroeconomic variables, and bootstrapped p-values ^a	14
1.4	Robust sign restrictions on impact based on the ACD model	18
2.1	Correlations and standard deviations (bandpass-filtered, 6-32 quarters)	43
2.2	The various facets of the MBC shock, variance contributions (VECM 1)	50
2.3	Average correlation of conditional time series (VECM1)	51
2.4	Variance contribution of the MBC shocks around the world (VECM1)	53
2.5	Inflation and the business-cycle (VECM1)	54
2.6	The MBC shock and capacity utilization, variance contributions (VECM 1)	56
2.7	MBC shock, variance contributions at long-run frequencies (80-∞ Q) (VECM 1)	57
2.8	MBC shock, variance contributions at long-run frequencies (80-∞ Q) (VECM 1)	58
2.9	Long-run shocks, contributions at long-run frequencies (80-∞ Q) (VECM 1) . .	59
2.10	Long-run shocks, contributions at business-cycle frequencies (VECM 1)	61
2.11	Variance contribution of the MBC shock around the world (VECM1)	65
2.12	Correlation of shocks (VECM1)	66
2.13	The various facets of the MBC shock, variance contributions (VECM1)	67
2.14	OECD EO dataset	72
2.15	Ohanian and Raffo dataset	73
2.16	Correlations and standard deviations (bandpass-filtered, 6-32 quarters)	75
2.17	Trade Openness Index: Trade in % of GDP	76
2.18	Exports in % of GDP	76
2.19	Imports in % of GDP	76
2.20	Correlations of conditional times series (VECM1)	79
2.21	The various facets of the MBC shock, variance contributions (VECM 1)	80

2.22	The various facets of the MBC shock, variance contributions (VAR)	82
2.23	Average correlation of conditional time series (VAR)	84
2.24	Long-run shocks, contributions at long-run frequencies ($80-\infty$ Q) (VAR)	84
2.25	Long-run shocks, contributions at business-cycle frequencies (VAR)	86
2.26	The various facets of the MBC shock, variance contributions (VECM 1)	87
2.27	The various facets of the MBC shock, variance contributions (VECM 1)	88
2.28	Correlation of shocks (VAR)	89
2.29	The various facets of the MBC shock, variance contributions (VAR)	90
3.1	Costs of leaning against real housing prices or real credit in terms of real GDP losses	109
3.2	Data Description	111
3.3	Data Transformation	111
3.4	Costs of leaning against real housing prices or real credit in terms of real GDP losses	114
4.1	Overview of experimental sessions and participants.	125
4.2	Summary descriptive statistics for delegates across all six experimental sessions.	126
4.3	Determinants of actual average CO2 emission reductions 2030-2050 grouped on experimental session level.	128

Introduction

This dissertation consists of four separate essays on macroeconomics. The first three chapters present empirical contributions that focus on quantifying the impact of various shocks using time series models frequently used in modern macroeconomics. The last chapter has a different methodological scope. It develops a novel experimental format to simulate international climate change negotiations. In the first chapter, I quantify the role of autonomous fluctuations in consumer confidence in driving business-cycle fluctuations in the United States (US). This chapter concludes that confidence shocks play only a minor role in US business-cycle fluctuations. In the second chapter, I identify the main driver of business-cycle fluctuations in six advanced economies to establish a common international pattern. I find that business-cycles in advanced economies are all alike and driven by the same type of aggregate demand shock. In the third chapter, co-authored with Larissa Schwaller, we identify monetary policy shocks to investigate the cost central banks incur when attempting to lean against housing prices or credit aggregates using their interest rate instrument. We find that the effect of such policies is limited and associated with significant real economic costs. In the fourth chapter, co-authored with Klaus Schmidt and Elisa Hofmann, we develop a novel experimental format to simulate international climate change negotiations. We compare negotiations on individual commitments, as in the Paris agreement, with negotiations on a common commitment, such as a uniform global carbon price. We find that there is a significant and positive impact on long-term emissions when negotiations focus on a uniform global carbon price.

Identifying and characterizing the causes and origins of business-cycle fluctuations is a long-standing quest of the macroeconomic profession. The existing literature has identified a plethora of shocks significant to business-cycle fluctuations. Prominent examples include technology shocks, such as neutral technology (N) shocks or investment-specific technology (IS) shocks, as well as monetary policy shocks, fiscal shocks, oil shocks, and uncertainty shocks. A comprehensive overview of commonly identified shocks is provided by Ramey (2016). In recent years, new literature has explored the idea that business-cycle fluctuations may originate predominantly from confidence shocks. Orthogonal to the fundamental state of the economy or news thereof, these shocks affect households' and firms' optimism or pessimism about the

short-term economic outlook, causing shifts in aggregate demand. In my first dissertation chapter, I show that several existing approaches to the identification of such confidence shocks are unreliable, as they identify a role for confidence shocks even when applied to models that do not feature them by construction. The reason that these approaches are unreliable is that they impose weak and generic restrictions that are easily satisfied by random rotations of the model's true structural disturbances. Working within a structural vector autoregression (SVAR) framework, I propose an alternative identification that imposes more stringent restrictions. The proposed scheme combines two parts. The first part identifies two permanent technology (N and IS) shocks with the 'max-share identification' approach pioneered by Uhlig (2003). In this approach, a shock is identified as the one that dominates the forecast error variance (FEV) of a particular variable at a specific horizon. Conditional on having identified permanent shocks, I use the approach of Arias, Rubio-Ramírez, et al. (2018) that imposes sign restrictions to identify several transitory shocks. I use the theoretical model of Angeletos et al. (2018), which includes confidence shocks in addition to traditional drivers of macroeconomic fluctuations, as a laboratory to test my identification scheme. The results indicate that the identification method always accurately recovers the permanent technology shocks. When applied to actual data for the US, the identification restrictions reveal that the two permanent technology shocks jointly explain a significant portion of all variables in the business-cycle band of 6-32 quarters. Based on the proposed identification scheme, the chapter concludes that confidence shocks only play a minor role in business-cycle fluctuations, while permanent technology shocks jointly explain a large portion of the observed variations. This conclusion is in line with Barsky and Sims (2012) and Fève and Guay (2019) who find a similarly minor role for autonomous fluctuations in confidence.

Angeletos et al. (2020) (ACD) recently introduced the concept of a 'main business-cycle' (MBC) shock to refer to the force that is responsible for the majority of business-cycle fluctuations in the key macroeconomic variables (output, investment, consumption, hours worked, and unemployment) in the US. Motivated by Lucas (1977), who famously observed that business-cycles are all alike, the objective of my second chapter is to investigate whether similar main drivers of business-cycles exist in the cross-section of advanced economies to assert that business-cycles are indeed all alike and driven by the same type of shock, the MBC shock. Estimating VARs and vector error correction models (VECMs), I identify the main drivers of business-cycle fluctuations in Australia, Canada, France, Italy, the United Kingdom, and the US to establish the existence of a common empirical template. Akin to the international business-cycle literature, I examine the comovement of the identified forces across countries and study the relationship between the main domestic business-cycle driver and variables related to international trade. To identify the main drivers of business-cycle fluctuations, I apply the

identification strategy of ACD. Their approach also builds on the max-share approach of Uhlig (2003). Three crucial differences stand out relative to the approach taken in the previous chapter. First, instead of maximizing the FEV of a particular variable at a specific long-run horizon in the time domain, spectral decomposition is used to isolate the variance of a particular variable within a specific frequency band. The use of spectral decomposition allows to examine different frequency components of the data in isolation. I consider two frequency bands, one associated with the business-cycle (6-32 quarters) and one associated with the long-run ($80-\infty$ quarters). Second, unlike the structural identification of transitory shocks via the sign restriction approach of Arias, Rubio-Ramírez, et al. (2018), ACD's approach does not impose any restrictions motivated by economic theory nor does it identify a structural shock as in the first chapter. Instead, it identifies a particular shock as the one that contributes the most to the volatility of a particular variable over a specific frequency band. The recovered shock should therefore be seen as the dominant driver of a specific variable at a particular frequency. Third, ACD's strategy consists of taking multiple cuts of the data by varying the targeted variable and frequency band. By targeting different variables, the same shock may be identified for multiple variables. Yet, these multidimensional cuts of the data help characterize the properties of the main business-cycle driver I aim to identify. Echoing Lucas (1977) in that business-cycles are all alike, I find support for the existence of an MBC shock, a single dominant force in each country that explains the bulk of macroeconomic fluctuations, triggers strong comovement in the main macroeconomic variables, and a countercyclical reaction of net exports as imports increase by more than exports. Moreover, each MBC shock is disconnected from inflation, the terms of trade, and the supply forces that drive economic activity in the long-run. In line with recent research in the international business-cycle literature (Levchenko and Pandalai-Nayar (2020) and Huo et al. (2020)) that propose demand-side shocks to be important drivers of international business-cycle comovement, I find that the identified MBCs play a significant role in explaining international output comovement.

Since the Great Financial Crisis (GFC), an intensive discussion around leaning against the wind (LAW) policies has emerged. LAW policies involve central banks pursuing a tighter monetary policy than what is consistent with inflation targeting, without taking financial stability effects into account. On the one hand, LAW policies can potentially reduce the probability and magnitude of crises. On the other hand, they can lead to a weaker economy, lower GDP, higher unemployment, and lower inflation rates. The dominant view among economists is that the costs of LAW policies outweigh the benefits by a significant amount (see Svensson (2017), Schularick et al. (2021), Benati (2021), and Benati (2022)). However, there is also a minority position that central banks should pursue financial stability-oriented monetary policy, which involves deviating temporarily from traditional objectives to avert a financial crisis when signs

of imbalances emerge (see BIS (2014) and BIS (2016)). In my third chapter, we measure the costs central banks face when implementing LAW policies using their interest rate instrument. We focus on leaning against housing prices and credit aggregates and evaluate the costs in terms of relative real GDP losses. We estimate VARs and use an external instrument identification strategy to identify monetary policy (MP) shocks. This approach has been developed by Stock and Watson (2012) and Mertens and Ravn (2013) and is commonly referred to as Proxy-SVAR. The advantage of using this framework is that it does not impose a particular theory on how a MP shock should affect any variables as done by using e.g., zero and sign restrictions. We use data for the US and consider the frequentist Proxy-SVAR approach by Gertler and Karadi (2015) and the Bayesian approach by Arias, Rubio-Ramírez, et al. (2021). As an external instrument, we use the orthogonalized monetary policy surprises (MPS) of Bauer and Swanson (2022). MPS are high-frequency changes in interest rates around Federal Open Market Committee (FOMC) announcements. Recently, some studies raised concern about the exogeneity of MPS because they found that they are systematically correlated with economic and financial data that is available prior to FOMC announcements. Bauer and Swanson (2022) correct for this issue by orthogonalizing their MPS with respect to economic and financial data that predate FOMC announcements. This treatment ensures the exogeneity of MPS and their validity as an instrument for identifying MP shocks. Based on our combination of state-of-the-art identification methods and the improved measures of MPS, we find that the impact of an unexpected 25 basis point interest rate hike on real housing prices or real credit is not significantly larger than the impact on real GDP. In line with the literature that advocates against LAW policies, we find that the ability of the Federal Reserve to correct disequilibria in asset price markets or credit markets appears limited and is associated with significant real economic costs.

Climate change can only be mitigated if sovereign countries negotiate effective international cooperation. The success of these negotiations depends on how they are structured. In the Paris agreement negotiations, countries negotiated a common, non-binding goal of how much to limit global warming. This goal is to be implemented by ‘nationally determined contributions’, i.e., individual commitments by the participating countries of how much to contribute to the common good. Weitzman (2014), Weitzman (2017a), Nordhaus (2015), Nordhaus (2019), MacKay et al. (2015), and others propose to focus climate negotiations on a uniform minimum price for carbon emissions instead. This negotiation design strives for a common commitment that builds on reciprocity. While there are several theoretical arguments for focusing climate change negotiations on a carbon price, almost no empirical or experimental evidence exists to test which negotiation protocol is more successful in achieving ambitious climate action. In my fourth chapter, we investigate whether negotiating a uniform commitment (such as a carbon price) achieves more ambitious CO_2 reductions than individual commitments (such as the ne-

gotiations that led to the Paris Agreement). We provide a novel field experiment that allows us to simulate climate change negotiations by the Conference of the Parties (COP), the main body of the United Nations Framework Convention on Climate Change (UNFCCC). We collaborated with Model United Nations (MUN) associations at Universities in Germany and Switzerland. MUN is a formally structured framework to simulate complex negotiations. Participants take on the role of diplomats and represent various countries to simulate actual committees of the UN at international conferences around the world. Delegates conduct research, adhere to formal rules of debate, negotiate in plenary and informal discussions, and vote on resolutions. We simulate two committees, each with the same ten countries, including major CO_2 emitters, major producers of fossil fuels, and developing countries. The two committees are experimental treatments that differ in only one respect. In one committee, individual commitments are negotiated, similar to the Paris Agreement. In the second committee, negotiations aim at a common uniform commitment with a uniform global carbon price. Our experiment is designed to simulate realistic, dynamic negotiations in a formally structured setting to create replicable conditions that enable the demonstration of statistically valid results. We find that focusing climate change negotiations on a single carbon price has a significant and positive impact on long-term emissions reductions, higher participation in the resolution, and more equal contributions than negotiating individual commitments as in the Paris agreement. Taken together, our results suggest that the international community should give the approach of negotiating a uniform global carbon price serious consideration.

Chapter 1

Identification of Confidence Shocks

Abstract

We make two contributions to the literature exploring the role of confidence shocks in macroeconomic fluctuations: *(i)* Working with the theoretical moving average (MA) representations of standard dynamic stochastic general equilibrium (DSGE) models, we show that several structural vector autoregression (SVAR) based approaches to the identification of confidence shocks are unreliable, as they identify such disturbances even when the model does not feature them. The problem is that the restrictions which are typically imposed are so weak and generic that they will always be satisfied with non-negligible probability by random rotations of the model's structural disturbances, irrespective of the fact that they do, or do not, include a pure confidence shock. *(ii)* We derive robust restrictions for the identification of confidence shocks based on the model of Angeletos et al. (2018), and working with the theoretical MA representation of the model, we evaluate their performance in identifying the shocks' true impulse response functions (IRFs) and fractions of forecast error variance (FEV). Our main finding is that our restrictions always allow us to exactly recover permanent investment-specific (IS) and neutral (N) technology shocks. When we impose our restrictions upon the data within a SVAR framework, we find that permanent IS and N shocks already jointly explain large-to-dominant portions of the FEV of all real macroeconomic variables in the business-cycle band of 6-32 quarters. In conclusion, we detect only a minor-to-negligible role for confidence shocks.

Acknowledgment: This is a revised and completed version of the paper titled 'Does Sentiment Matter?' that was started with Luca Benati, whom I would like to thank for his continuous support and express my deepest gratitude for his invaluable advice from start to finish of this project. Moreover, I am thankful for useful comments and suggestions from Fabrice Collard, Harris Dellas, Annina Heini, the Macro Group seminar participants in Bern, the participants of the 2018 Theories and Methods in Macroeconomics in Paris, the 2018 Young Swiss Economists Meeting in Zurich, the 2017 North American Summer Meeting of the Econometric Society in St. Louis, and the 2017 Asian Meeting of the Econometric Society in Hong Kong. Special thanks to Allen Li, of the Conference Board, for kindly providing the Board's data on US consumer and CEO confidence.

1.1 Introduction

In recent years, a vast literature has explored the idea that macroeconomic fluctuations may originate largely or predominantly from autonomous fluctuations in consumer confidence or sentiment.¹ In this paper, we contribute to this literature by making two main points:

First, working with the theoretical moving average (MA) representations of several standard dynamic stochastic general equilibrium (DSGE) models, we show that existing approaches to the identification of ‘sentiment shocks’ are unreliable, as they identify such disturbances even when the model does not feature them. The approach used by Beaudry, Nam, et al. (2011) and Nam and Wang (2019), for example, identifies sentiment shocks within Smets and Wouters (2007)’s (henceforth SW) model, which, by construction, does not feature them. By the same token, we show that the popular approach of identifying sentiment shocks as innovations to measures of consumer confidence detects an entirely spurious role for autonomous fluctuations in sentiment even when the DSGE model is block-recursive by construction so that consumer confidence does not Granger-cause other variables, and ‘pure sentiment’ shocks only affect consumer confidence. We argue that this is not by chance, but that it should be expected since the restrictions imposed by the likes of Beaudry, Nam, et al. (2011), Nam and Wang (2019), and Déés and Güntner (2014) are extremely weak and generic and will therefore be satisfied with a non-negligible probability by random rotations of the model’s true structural disturbances, irrespective of the fact that such disturbances do or do not include a pure sentiment shock.

Second, working within a structural vector autoregression (SVAR) framework, we derive robust sign restrictions - along the lines of Canova and Paustian (2011) - from the theoretical DSGE model of Angeletos et al. (2018) (henceforth ACD). Their model belongs to a novel class of theoretical models that introduces waves of optimism and pessimism about the short-run outlook of the economy. In their model, a positive innovation in optimism (a sentiment shock) causes agents to have optimistic expectations of income and profitability in the short-run, ultimately triggering a short-lived macroeconomic boom in main macroeconomic aggregates. Crucially, the sentiment shock is independent of technological fundamentals such as total factor productivity (TFP) or news thereof. The benefit of using ACD’s model is that it introduces a tractable role for sentiment shocks about the short-run economic outlook with only a minimal increase in the state space.² Additionally, ACD’s model contains the familiar bells and whistles of the New Keynesian DSGE literature and features pure sentiment shocks alongside more

¹We use the terms beliefs, confidence, and sentiment interchangeably in what follows. Moreover, we understand a sentiment shock as a shock that causes agents to become more optimistic about the short-term economic outlook in a way that is orthogonal to both the fundamental (i.e., TFP) and news about the fundamental - but nevertheless exhibits a causal role on the economy.

²This is achieved by introducing heterogeneous priors and thus departing from the standard common prior assumption. Heterogeneous priors allow accommodating variation in economic activity due to shifts in expectations. In simple terms, these allow to introduce shocks to the forecasts of forecasts of others. On the other hand, with common priors, forecasting the forecasts of others is trivial as they are all the same.

traditional drivers of macroeconomic fluctuations, such as investment-specific (IS) and neutral (N) permanent technology shocks and (transitory) monetary policy shocks. These features allow us to derive robust sign restrictions via the methodology of Canova and Paustian (2011) for several shock-variable pairs, which also have well-defined empirical counterparts. To foreshadow, our identifying restrictions combine two parts. The first part relies on the ‘maximum FEV’ approach pioneered by Uhlig (2003) and Uhlig (2004), in which a shock is identified as the one that dominates the forecast error variance (FEV) of a particular variable at a specific horizon.³ The second part of our identifying restrictions relies on imposing robust zero and sign restrictions as proposed by Arias, Rubio-Ramírez, et al. (2018). Working with the theoretical MA representation of the ACD model, we then evaluate the performance of our restrictions via several experiments. Our restrictions always allow us to recover permanent IS and N shocks exactly. Regarding the remaining structural disturbances, we find, on the one hand, that (i) when IS shocks explain only a negligible fraction of the FEV of the variables considered and (ii) the confidence shock explains a large portion, then our methodology underestimates the IRF and FEV of the confidence shock, but still recovers all other shocks precisely. This result holds regardless of the presence of uncertainty about the random rotation matrices we use to impose robust sign restrictions. On the other hand, we find that when the importance of IS and N shocks is closer to what we later observe in the empirical exercise, where IS and N shocks explain large-to-dominant portions of the FEV of all variables, then the identification of all shocks - including the confidence shock - works very well. The only slight qualification to this result is consumption. We impose our identifying restrictions on actual data for the United States (US) and work with systems featuring standard macroeconomic variables and several alternative measures of consumer confidence. An extremely robust finding therein is that permanent IS and N shocks already jointly explain large-to-dominant portions of the FEV of all variables at the business-cycle band of 6-32 quarters. Since (i) the presence of these two disturbances is generally well-established in the macroeconomic profession;⁴ (ii) the way we identify them, via Uhlig (2003) and Uhlig (2004)’s maximum FEV approach is standard; and (iii), as mentioned earlier, our restrictions allow the exact recovery of the two shocks in population conditional on the theoretical MA representation of ACD’s model, the fact that these two shocks jointly explain large-to-dominant portions of the FEV of all variables at business-cycle frequencies puts a robust upper bound on the role that pure sentiment shocks might play. To put it differently,

³An alternative version of this identification, which we do not explore in this paper, is to use spectral decomposition to target the variance of a particular variable within a particular frequency band. This identification is, for example, applied by Angeletos et al. (2020) and will be introduced in the next chapter of this thesis. Moreover, note that the fraction of FEV explained by a shock is a common measure to quantify the relative importance of a shock in explaining the variation of a macroeconomic variable.

⁴For IS shocks, see, e.g., Greenwood, Hercowitz, and Huffman (1988), Greenwood, Hercowitz, and Krusell (1997), Greenwood, Hercowitz, and Krusell (2000), and Fisher (2006). For N shocks, see e.g. Barsky and Sims (2011).

we find that permanent IS and N shocks already play a large part in driving business-cycle fluctuations. Hence, even in the circumstance in which sentiment shocks explained all of the residual FEV of macroeconomic variables not explained by IS and N shocks, this would not amount to much. A key point to stress is that this result is independent of the fact that the other shocks are identified via the robust sign restrictions the ACD model implies. Once we consider that beyond IS and N shocks, several other macroeconomic disturbances (e.g., monetary and fiscal policy shocks) are unquestionably present, and that they do in fact account for some of the FEV of the data, the possibility that sentiment shocks might play a dominant role in driving business-cycle fluctuations appears unlikely. This observation is corroborated by the fact that our identified sentiment shocks consistently explain small-to-negligible fractions of the FEV of all variables, including all of the confidence indices themselves. Our results are therefore consistent with those of Fève and Guay (2019, p. 877), who, while working within an SVAR framework, and identifying restrictions different from ours, conclude that sentiments shocks ‘explain little of quantities and prices’, and they ‘mostly appear as an idiosyncratic component of confidence.’

Finally, in passing, we reconsider the issue of whether survey measures of either consumer or CEO confidence Granger-cause standard macroeconomic time series. We consider eleven indices of consumer confidence from the Michigan survey and three from the Conference Board, as well as four indices of CEO confidence from the Conference Board. The null of no Granger-causality of confidence indices onto macroeconomic time series is almost uniformly rejected, most of the time very strongly. This is the case not only for small systems featuring just two additional macro variables - for which this result should logically be expected - but also for larger systems featuring six or twelve additional macro series. The rejection of no Granger-causality is compatible with the notion that the considered confidence indices do indeed contain independent information which is not contained in the other macroeconomic series considered and might be pure sentiment. This finding motivates our choice to include a measure of sentiment in our empirical application.

The remainder of this paper is organized as follows: in Section 1.2, we show that several existing approaches to identifying sentiment shocks are unreliable, as they identify such disturbances even when the model does not feature them. In Section 1.3, we consider the issue of whether confidence indices Granger-cause standard macroeconomic time series. In Section 1.4, we derive robust sign restrictions from the ACD model and, working with the theoretical MA representation of the model, investigate whether these restrictions allow for the recovery of the structural disturbances. In Section 1.5, we estimate SVARs featuring standard macroeconomic variables and several alternative measures of consumer confidence and impose our identifying restrictions. We conclude in Section 1.6.

1.2 Assessing Existing Approaches to Identifying Confidence Shocks

1.2.1 Beaudry et al. (2011) and Nam and Wang (2016)

Working within a SVAR framework, Beaudry, Nam, et al. (2011) and Nam and Wang (2019) identify sentiment shocks by imposing the set of zero and sign restrictions on impact, which we report in the table below, where ‘0’, and ‘+’ mean that the impact has been restricted to be zero and non-negative, respectively, whereas ‘?’ means that it has been left unrestricted. Finally, in all of the models they estimate, GDP, investment, and hours are left unrestricted.

Table 1.1: Zero and sign restrictions from Beaudry, Nam, et al. (2011) and Nam and Wang (2019)

	TFP	Stock price	Consumption	Real interest rate
Scheme I	0	+	?	?
Scheme II	0	+	+	?
Scheme III	0	+	+	+

Beaudry, Nam, et al. (2011) comment on the restrictions:

‘The impulse responses of variables are restricted to be zero (0) on impact, non-negative (+) on impact, or unrestricted (blank) in either the five-variable system (with TFP, Stock Price, Consumption, Real Interest Rate, Hours), the seven-variable system (with TFP, Stock Price, Consumption, Real Interest Rate, Hours, Investment, Output), or the eight-variable systems where an additional variable of interest is added to the seven-variable system.’

Based on the theoretical MA representation of SW’s model - i.e., a model which does not feature pure sentiment shocks by construction - we now show that the mechanical application of these restrictions spuriously identifies sentiment shocks which turn out to explain non-negligible fractions of the FEV of macroeconomic variables. An important point to stress is that we perform our analysis while working in population. Therefore, our result does not hinge on small sample issues or the like. Instead, it simply stems from the fact that the restrictions reported in Table 1.1 are so weak that they are inevitably going to be satisfied by a non-negligible fraction of (random) rotations of an initial estimate of the model’s structural impact matrix. Another way of putting this is that these restrictions are so generic that, if a researcher were to randomly rotate the Cholesky factor of the covariance matrix of innovations of SW’s model,

they will obtain, with a non-negligible probability, a linear combination of the model’s true structural shocks that end up satisfying these restrictions. These identified ‘sentiment’ shocks will however be nothing but linear combinations of the seven true structural shocks perturbing SW’s model, which just happen to satisfy the weak restrictions reported in the Table 1.1. With regard to the above restrictions, we uniquely focus on Scheme III for two reasons. First, out of the three Schemes in Table 1.1, it imposes the largest number of restrictions. Based on the above argumentation of weak identification, it may therefore be regarded as more reliable than the other schemes. Second, we rule out Scheme I as it is manifestly problematic by itself for a very simple reason. As in Beaudry and Portier (2006), a shock which leaves TFP unaffected on impact, and causes stock prices to rise, could well be a TFP news shock. In fact, this is also actually one of the schemes used by Beaudry and Portier (2006) in order to identify such shocks. This hence provides an extreme example of a problem we will discuss in-depth in this section: Existing approaches to the identification of sentiment shocks suffer from the shortcoming that they tend to incorrectly identify other disturbances as sentiment shocks.

Let us now turn to the application of Scheme III in population. To this end, we solve SW’s model conditional on their modal estimates for either the 1966:Q1-1979:Q2 (‘Great inflation’ period) or the 1984:Q1-2004:Q4 (‘Great Moderation’ period) subsample. Based on the state-space representation of the model, we then compute the structural MA representation for the following seven variables, which we collect in the vector Y_t : Neutral technology (TFP), GDP, consumption, investment, the FED Funds rate, inflation, and Tobin’s Q (which is conceptually related to stock prices):

$$Y_t = A_0\epsilon_t + A_1\epsilon_{t-1} + A_2\epsilon_{t-2} + A_3\epsilon_{t-3} + A_4\epsilon_{t-4} + \dots \quad (1.2.1)$$

In the above equation, the A_j ’s are the MA matrices in the structural MA representation of SW’s model and ϵ_t is the vector collecting the seven structural shocks in SW’s model (neutral technology, investment-specific technology, risk premium, exogenous spending, price mark-up, wage mark-up, and monetary policy). Further, note that in the above Equation (1.2.1), A_0 is the true structural impact matrix of the shocks for SW’s model. To identify the structural impact matrix based on the restrictions in Scheme III as reported in Table 1.1, we then proceed as follows: First, we take as the initial estimate of the structural impact matrix the Cholesky factor of A_0A_0' . Let this starting estimate be A_0^* . Thereafter, we rotate A_0^* via the algorithm for combining zero and sign restrictions proposed by Arias, Rubio-Ramírez, et al. (2018). For $j = 1, 2, 3, \dots, N$, with $N = 100,000$, we consider a single random rotation matrix implementing the zero restrictions on A_0^j (with A_0^j being the candidate structural impact matrix), which we generate via Arias, Rubio-Ramírez, et al. (2014)’s Algorithm 5.⁵ Then, if the sign restrictions

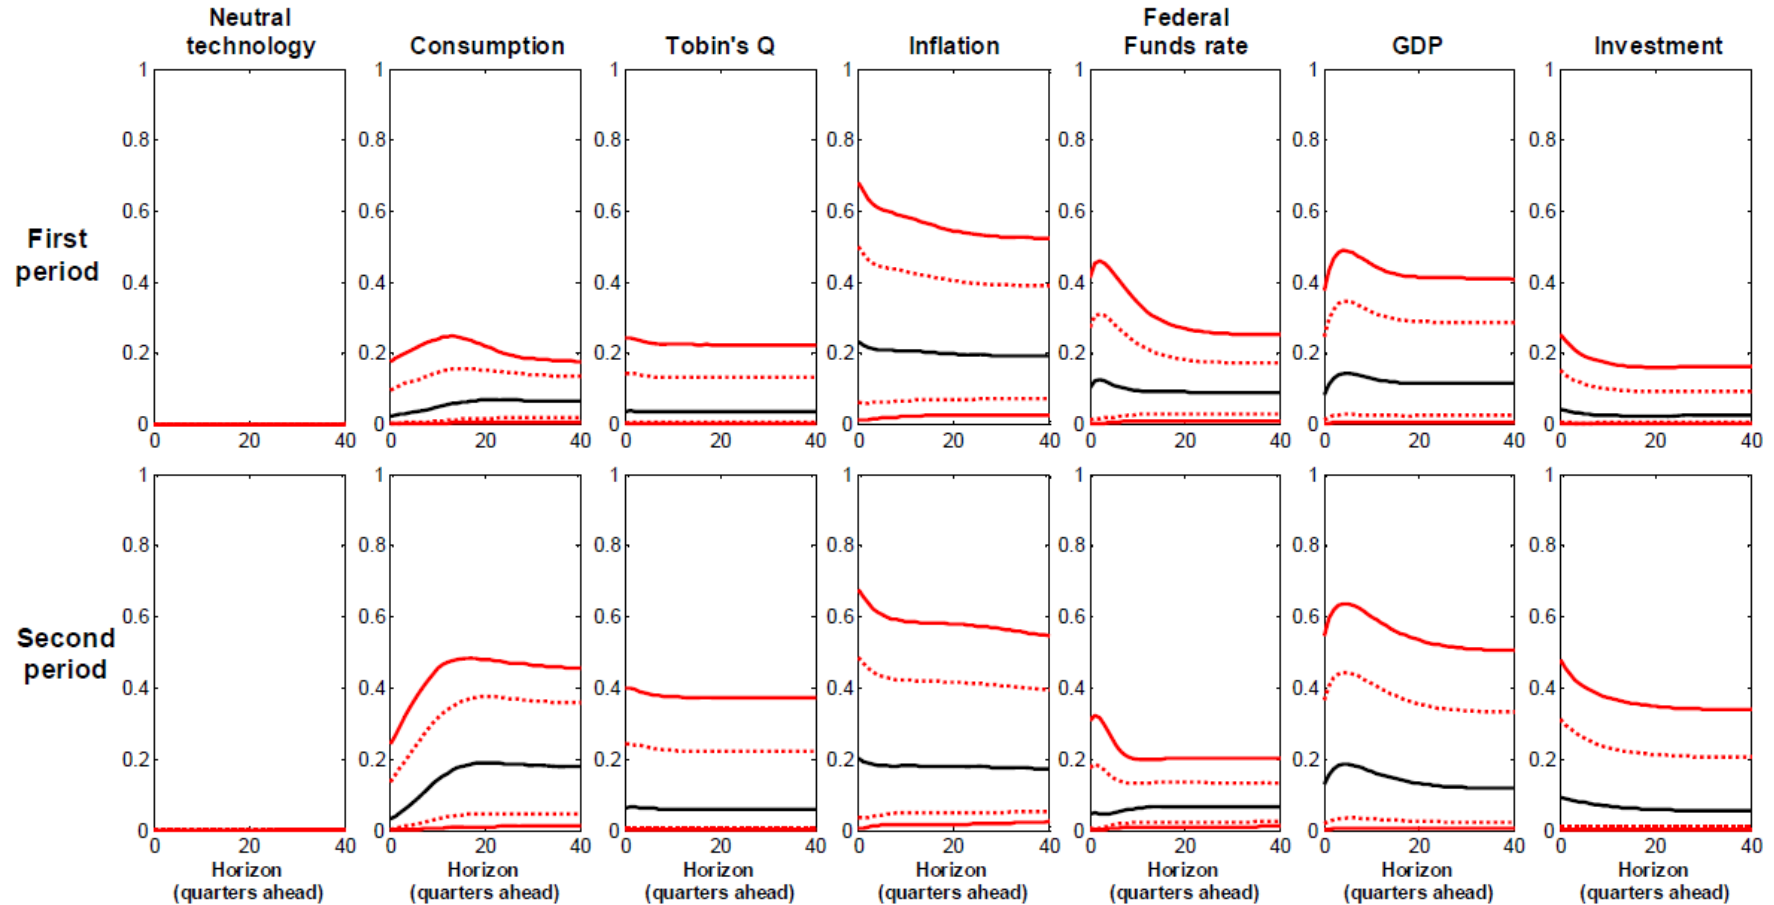
⁵See Arias, Rubio-Ramírez, et al. (2014) p. 18 (the version of the paper we are referring to is dated September

are all satisfied, we keep A_0^j . Otherwise, we discard it. In this way, we build up the statistical distribution of the A_0^j 's conditional on (i) SW's model's covariance matrix of innovations, and (ii) the zero and sign restrictions reported in Table 1.1. Finally, for each of the A_0^j 's, we compute the fraction of FEV of the seven variables at horizons up to 40 quarters which is explained by the identified sentiment shock.

Figure 1.1 reports the median, together with the 16-84, and the 5-95 percentiles of the distributions of the fractions of FEV of either variable explained by the identified sentiment shock. Based on the medians of the distributions, we see that the fractions of explained FEV in the case of TFP are exactly zero at all horizons. This reflects the facts that (i) in the VAR representation of SW's model, TFP only depends on itself lagged, and (ii) the identified sentiment shocks, by assumption, do not impact TFP at $t=0$. With regard to the other variables, and if we focus on the 5-95 credible sets, we see that the identified sentiment shock explains non-negligible fractions of FEV. This is the case, in particular, for inflation, GDP, and consumption in the second period. The objective truth that the SW model does not contain pure sentiment shocks by construction illustrates that the obtained results are entirely spurious and the figment of imposing upon the data generating process (DGP) a set of very weak restrictions. The above experiment illustrates the aforementioned argument that the mechanical application of weak restrictions allows for the spurious identification of sentiment shocks, even if they are not present in the underlying macroeconomic model.

7, 2014). Notice that although they discuss a more general algorithm based on Gibbs-sampling, as they point out '[...] when the researcher is interested in identifying only one shock, the Gibbs sample step in Algorithm 4 is not necessary, and one should obtain q from Algorithm 5.'

Figure 1.1: Spuriously identifying confidence shocks in SW's model via Beaudry et al.'s (2011) identification strategy: fractions of FEV explained by the identified confidence shock



Black line: median of the distribution, red line: 5-95 percentiles of the distribution, red dotted line: 16-84 percentile of the distribution. The first period refers to the 1966:Q1-1979:Q2 sample. The second period refers to the 1984:Q1-2004:Q4 sample.

1.2.2 Identifying shocks to measures of consumer sentiment

Another approach to identifying sentiment shocks is based on the notion of identifying innovations to measures of consumer confidence (see, e.g., Déés and Güntner (2014)). As we now show by example, this approach is also unreliable for the simple reason that it mechanically identifies sentiment shocks even within environments in which consumer sentiment is, by construction, a pure linear measure of macroeconomic variables.

The easiest way to illustrate this is via the standard three-equations backward- and forward-looking New Keynesian model,

$$R_t = \rho R_{t-1} + (1 - \rho)[\phi_\pi \pi_t + \phi_y y_t] + \epsilon_{R,t} \quad (1.2.2)$$

$$\pi_t = \frac{\beta}{1 + \alpha\beta} \pi_{t+1|t} + \frac{\alpha}{1 + \alpha\beta} \pi_{t-1} + \kappa y_t + u_t \quad (1.2.3)$$

$$y_t = \gamma y_{t+1|t} + (1 - \gamma)y_{t-1} - \sigma^{-1}[R_t - \pi_{t+1|t}] + v_t \quad (1.2.4)$$

where the notation is as follows: R_t , π_t , and y_t are the nominal interest rate, inflation, and the output gap, respectively. $\epsilon_{R,t}$, u_t , and v_t are three white noise structural disturbances, $\epsilon_{R,t} \sim N(0, \sigma_R^2)$, $u_t \sim N(0, \sigma_u^2)$, $v_t \sim N(0, \sigma_v^2)$ to the monetary policy rule, the New Keynesian Phillips curve, and the dynamic IS curve, respectively. We calibrate this model based on Benati (2008)'s Table 12 modal Bayesian estimates for the US for the full sample, and for the period after the Volcker stabilization, respectively. Conditional on these parameters' values, the model (1.2.2)-(1.2.4) has a structural VAR(2) representation. It is to be noted that in this model, output is equal to consumption, i.e. $y_t = c_t$. In addition to the model as outlined above, we artificially create a measure of consumer confidence, Ξ_t , such that $\Xi_t = c_t + \eta_t$, $\eta_t \sim N(0, \sigma_\eta^2)$. In this case, η_t is an authentic, exogenous innovation to consumer confidence, which we make as small as possible. The only constraint is that the resulting VAR representation for $Y_t \equiv [R_t, \pi_t, y_t, \Xi_t]'$ is not stochastically singular. We therefore set $\sigma_\eta^2 = 10^{-12}$. This ensures that the true fractions of FEV of R_t , π_t , and y_t explained by exogenous innovations to consumer confidence (η_t) are, for all practical purposes, nil. The resulting structural VAR(2) representation for Y_t for the full sample period is then given by (1.2.5) below. Note first that up to η_t , the fourth row of (1.2.5) is identical to the third, thus reflecting the fact that $\Xi_t = c_t + \eta_t$. Second, in this model economy, by construction, the measure of consumer confidence (Ξ_t) is Granger-caused by all of the other variables in the system, but it does not Granger-cause any of them. This means that within (1.2.5), the role played by authentic, exogenous shocks to consumer confidence in explaining fluctuations in other macroeconomic aggregates is virtually nil.

We now apply to (1.2.5) the methodology of Déés and Güntner (2014), which identifies sentiment shocks based on the restriction that, in response to a positive 'wave of optimism', consumer confidence, consumption, and the real interest rate do not decrease. As before, we

work in population, that is, based on the theoretical structural VAR representation (1.2.5).

$$\begin{bmatrix} R_t \\ \pi_t \\ c_t \\ \Xi_t \end{bmatrix} \equiv Y_t = \underbrace{\begin{bmatrix} 1.1490 & 0.2562 & 0.5352 & 0 \\ -0.0258 & 0.8269 & 0.0920 & 0 \\ 0.0401 & 0.2089 & 1.2940 & 0 \\ 0.0401 & 0.2089 & 1.2940 & 0 \end{bmatrix}}_{B_1} Y_{t-1} + \underbrace{\begin{bmatrix} -0.2362 & -0.1569 & -0.3476 & 0 \\ -0.0016 & -0.0022 & -0.0350 & 0 \\ -0.0711 & -0.1705 & -0.3521 & 0 \\ -0.0711 & -0.1705 & -0.3521 & 0 \end{bmatrix}}_{B_2} Y_{t-2} + \underbrace{\begin{bmatrix} 0.5789 & -0.1029 & -0.3515 & 0 \\ -0.0156 & -0.7151 & 0.4581 & 0 \\ -0.3053 & 0.3859 & 0.3555 & 0 \\ -0.3053 & 0.3859 & 0.3555 & 10^{-6} \end{bmatrix}}_{A_0} \epsilon_t, \quad (1.2.5)$$

We proceed by taking the matrix A_0 in (1.2.5) as the starting estimate of the structural impact matrix we want to identify via sign restrictions. We generate $N = 100,000$ random rotation matrices via the algorithm proposed by Déés and Güntner (2014). For each random rotation matrix Q_j , $j = 1, 2, \dots, N$, we then compute the corresponding candidate impact matrix $A_{0,j}^* = A_0 Q_j$. Finally, we keep, among all of the $A_{0,j}^*$'s, only those satisfying the previously mentioned sign restrictions, thus obtaining their distribution.

Figure 1.2 shows the results. Although the median estimates are uniformly quite low, once again, it must be stressed that these results are entirely spurious. The objective truth in this model economy is that for all series, the role played by exogenous innovations to consumer sentiment is uniformly nil by construction. Further, (i) if we focus instead on the 5-95 credible sets, the results are uniformly bad, with the upper bound for inflation being around 80 percent for either period; and (ii) in a few cases - inflation for the full sample and the other series for the period after the Volcker stabilization - even the median fractions of FEV are clearly non-negligible. We hence observe that the issue of weak identification also persists in cases where the identification is additionally based on identifying innovations to measures of consumer confidence.

1.2.3 Why do these approaches fail?

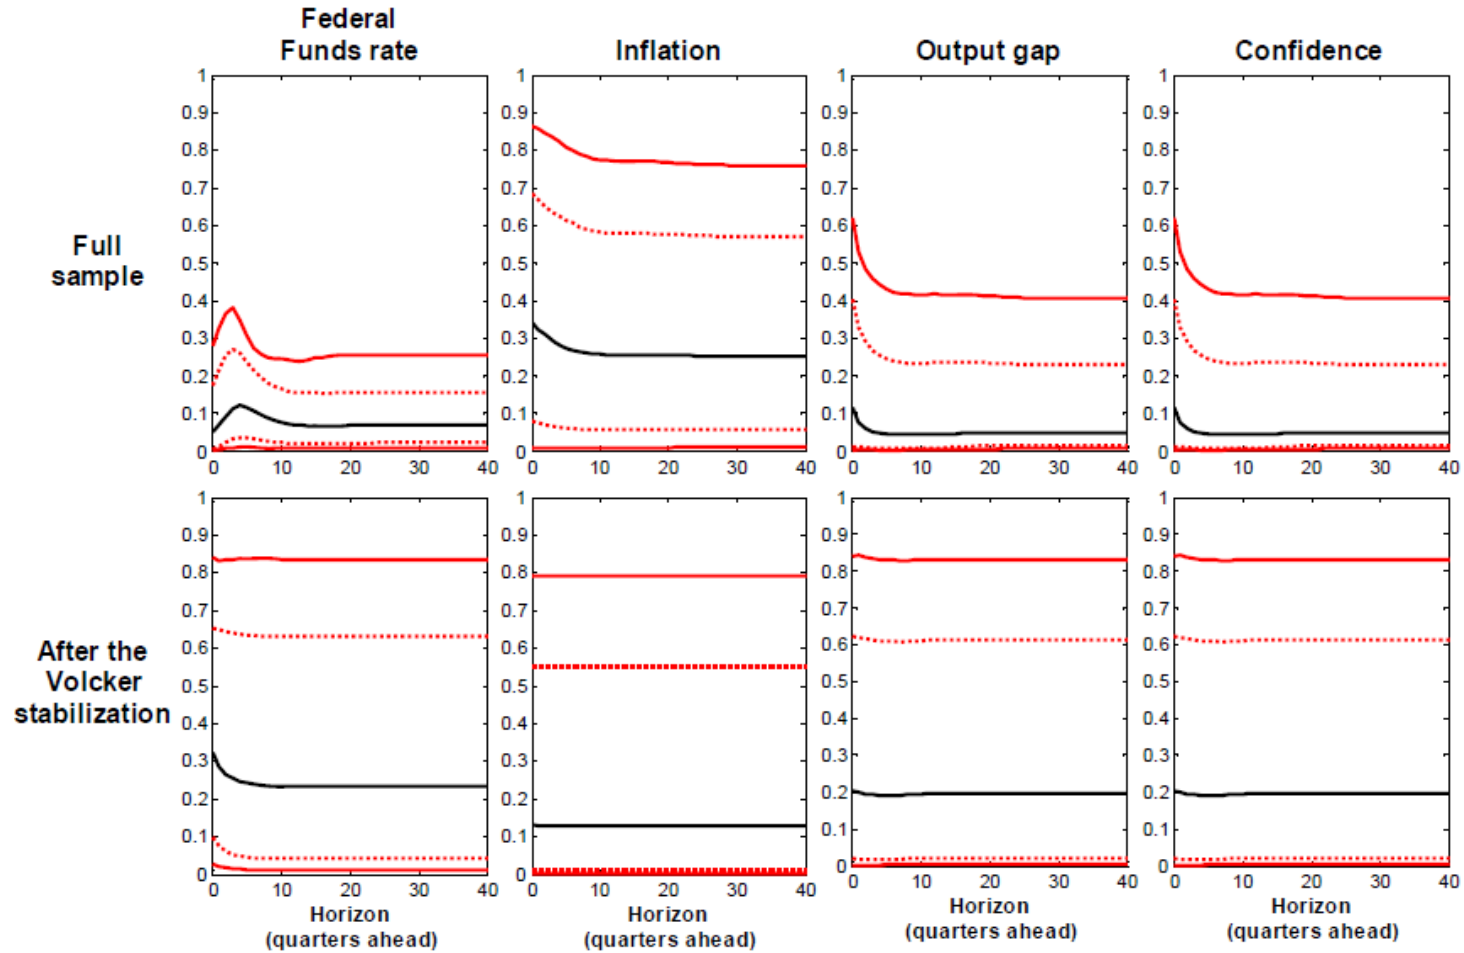
Why do the approaches discussed in the previous sub-sections fail to correctly identify the non-existent role of pure sentiment shocks within the underlying DGPs? The reason is straightforward, in that the restrictions imposed by Beaudry, Nam, et al. (2011), Nam and Wang (2019), and Déés and Güntner (2014) are so weak and generic that they will necessarily be satisfied with

a non-negligible probability by random rotations of the model's true structural disturbances, irrespective of the fact that such disturbances do or do not include a pure sentiment shock.

Conceptually in line with Canova and Paustian (2011) - and more generally with Fry and Pagan (2011)'s critique of sign restrictions as being 'weak information' - in what follows, we will impose upon our VARs a significantly more informative set of restrictions, which will allow us to obtain sharper inference. All of our restrictions will be derived based on the model of ACD. To anticipate, those pertaining to permanent IS and N shocks are standard in the literature (e.g., IS shocks are assumed to be the only driver of the unit root component of the relative price of investment (RPI)), whereas those pertaining to the other (transitory) disturbances will be derived as robust sign restrictions as in Canova and Paustian (2011).

Before doing so, we reconsider the issue of whether survey measures of consumer and CEO confidence Granger-cause standard macroeconomic time series. A key reason for doing so is that if, in fact, these measures did Granger-cause macroeconomic variables, this would suggest that such indices do indeed contain information about sentiment which is not contained in other macroeconomic series, thus strengthening the case for including them in the VARs.

Figure 1.2: Spuriously identifying confidence shocks in the standard New Keynesian model via sign restrictions: fractions of FEV explained by the identified confidence shock



Black line: Median of distribution, red line: 5-95 percentiles of the distribution, red dotted line: 16-84 percentile of the distribution.

1.3 Do Survey Measures of Confidence Granger-Cause Macro Variables?

In this section we perform Granger-causality tests of either consumer or CEO confidence indices onto macroeconomic variables. As we will see, the null of no Granger-causality of survey measures onto macro variables is strongly and near-uniformly rejected across the board, even for systems featuring either six or twelve additional variables. Although this is not hard proof that such additional information in measures of confidence is pure sentiment, it is at the very least compatible with such a notion and strengthens the case for including such measures in our empirical exercise in the following sections.

1.3.1 Testing for no Granger-causality

Tables 1.2 and 1.3 report Wald tests for the null hypothesis of no Granger-causality of confidence indices onto macroeconomic variables, together with bootstrapped p-values. The p-values (which are based on 10,000 bootstrap replications for each model) have been computed by bootstrapping the VAR model estimated under the null, that is, by imposing that confidence indices do not Granger-cause the other variables in the VAR.

The model estimated under the null has then been bootstrapped, and based on each bootstrapped replication we have performed the same Wald test for no Granger-causality we had previously performed based on the actual data, thus building up the empirical distribution of the test statistic under the null.

For either confidence index - either from the University of Michigan, or from the Conference Board; and either pertaining to consumers or to CEOs - we consider two sample periods. The first sample excludes the Zero Lower Bound (ZLB) period, which means that we end it in 2008:Q3, the last quarter for which we can reasonably assume that the ZLB was not binding. The second sample also includes the period since 2008:Q3.

For either period and confidence index, we consider three different models, which feature 3, 7, or 13 variables, respectively. The smallest model includes, beyond the confidence index, the log-difference of real GDP per capita and the logarithm of the consumption/GDP ratio. The intermediate model also includes the GDP deflator inflation, the logarithm of hours worked per capita, the log-difference of the RPI, and the ex post real FED Funds rate. Finally, the largest model further includes the logarithm of the investment/GDP ratio, the BAA-AAA spread, the spread between the 1-year government bond yield and the FED Funds rate; and the spread between the 3-year and the 1-year, the 5-year and the 3-year, and the 10-year and the 5-year government bond yields. All of the series included in the VARs are $I(0)$ according to the results of Elliott et al. (1996) unit root tests.

Table 1.2: Wald tests for Granger-causality of consumer confidence indices from the Conference Board onto macroeconomic variables, and bootstrapped p-values^a

	<i>Excluding the ZLB</i>			<i>Full samples</i>		
	<i>N=3</i>	<i>N=7</i>	<i>N=13</i>	<i>N=3</i>	<i>N=7</i>	<i>N=13</i>
<i>Consumer Confidence Index</i>						
Index 1: Overall Index	1.79 (0.03)	5.38 (0.01)	8.92 (0.05)	1.74 (0.02)	6.05 (0.00)	8.40 (0.01)
Index 2: Present Situation	2.39 (0.00)	6.79 (0.00)	8.72 (0.07)	1.72 (0.02)	5.91 (0.00)	
Index 3: Expectations	1.60 (0.04)	7.60 (0.04)	8.14 (0.09)	1.76 (0.01)	4.66 (0.00)	8.79 (0.00)
<i>Measure of CEO Confidence™</i>						
Index 4: Overall index	1.90 (0.02)	4.99 (0.01)	7.44 (0.11)	1.91 (0.01)	4.40 (0.00)	7.061 (0.03)
Index 5: Current Economic Conditions vs. 6 Months Ago	1.89 (0.01)	4.00 (0.03)	4.65 (0.78)	2.04 (0.00)	3.41 (0.03)	5.92 (0.12)
Index 6: Expectations for Economy, 6 Months Ahead	1.51 (0.05)	5.40 (0.00)	8.76 (0.03)	1.32 (0.05)	4.95 (0.00)	7.92 (0.01)
Index 7: Expectations for Own Industry 6 Months Ahead	1.85 (0.02)	5.60 (0.00)	11.31 (0.00)	1.97 (0.01)	5.23 (0.00)	7.41(0.02)

^a Based on 10,000 bootstrap replications.

Table 1.3: Wald tests for Granger-causality of consumer confidence indices from the Michigan survey onto macroeconomic variables, and bootstrapped p-values^a

	<i>Excluding the ZLB</i>			<i>Full samples</i>		
	<i>N=3</i>	<i>N=7</i>	<i>N=13</i>	<i>N=3</i>	<i>N=7</i>	<i>N=13</i>
<i>Consumer Confidence Index</i>						
Index 1: Overall Index	2.28 (0.00)	3.73 (0.01)	5.36 (0.10)	2.50 (0.00)	3.65 (0.00)	5.88 (0.01)
Index 2: Current Index	1.91 (0.00)	3.45 (0.01)	5.15 (0.14)	2.22 (0.00)	4.00 (0.00)	6.97 (0.00)
Index 3: Expected Index	2.12 (0.00)	3.74 (0.01)	5.89 (0.04)	2.15 (0.00)	3.18 (0.01)	5.43 (0.03)
<i>Index Components</i>						
Index 4: Personal Finances: Expected	1.62 (0.01)	2.99 (0.05)	4.49 (0.31)	1.40 (0.01)	2.57 (0.07)	5.08 (0.05)
Index 5: Business Conditions: 12 Months Ahead	2.20 (0.00)	3.98 (0.00)	5.90 (0.04)	2.33 (0.00)	3.52 (0.00)	5.48 (0.03)
Index 6: Business Conditions: 5 Years Ahead	1.86 (0.00)	3.52 (0.01)	6.00 (0.04)	1.78 (0.00)	2.81 (0.04)	4.81 (0.10)
<i>Answers to Questions^b</i>						
Index 7: Question 1	1.62 (0.01)	3.03 (0.04)	4.56 (0.28)	1.45 (0.01)	2.67 (0.06)	5.25 (0.040)
Index 8: Question 2	1.24 (0.05)	2.24 (0.33)	5.36 (0.18)	1.44 (0.01)	2.15 (0.26)	5.48 (0.05)
Index 9: Question 3	2.38 (0.01)	3.86 (0.01)	6.59 (0.03)	2.06 (0.00)	3.53 (0.01)	5.33 (0.07)
Index 10: Question 4	2.20 (0.00)	3.98 (0.00)	5.90 (0.04)	2.33 (0.00)	3.52 (0.00)	5.48 (0.03)
Index 11: Question 5	1.86 (0.00)	3.52 (0.01)	6.00 (0.03)	1.78 (0.00)	2.81 (0.04)	4.82 (0.10)

^a Based on 10,000 bootstrap replications. ^b See respective Questions for each Index in Appendix A.

Evidence from the two tables rejects almost uniformly and typically very strongly, the null hypothesis of no Granger-causality of confidence indices for the other variables included in the VARs. Notably, this is the case not only for systems featuring either 3 or 7 variables - for which a possible explanation could be that confidence indices do not contain any idiosyncratic information but are instead just ‘proxying’ for macro variables which have been left out from the VAR - but also for the largest systems. To be sure, it could well be the case that if we had worked with (e.g.) factor-augmented vector autoregressive (FAVAR) models, including factors

extracted from large panels of macro series, the null of no Granger-causality might not have been rejected. As they stand, however, the results in Tables 1.2 and 1.3 are at the very least compatible with the notion that confidence indices do indeed contain idiosyncratic information not contained in the other macroeconomic series considered, and might well be pure sentiment. This motivates our choice, in the rest of the paper, to include some confidence indices in the VARs, and to impose the restriction that pure sentiment shocks do not leave consumer or CEO confidence indices unchanged on impact (i.e., within the quarter).

1.4 Methodology

In this section, we now describe our own methodology for the identification of sentiment shocks. We start by introducing our empirical strategy for actual US data. This is then followed by a description of how we derive robust sign restriction from the model of ACD and how our identification fares in recovering the true structural disturbances in population.

1.4.1 Model and Variables

In what follows we will work with the VAR(p) model

$$Y_t = B_0 + B_1 Y_{t-1} + \dots + B_p Y_{t-p} + u_t, \quad E[u_t u_t'] = \Omega \quad (1.4.1)$$

where Y_t features (in this order) the logarithm of the RPI; the logarithm of real chain-weighted consumption of non-durables and services, GDP, and hours worked per capita; GDP deflator inflation; the Federal Funds rate; and either a consumer or a CEO confidence index (from either the University of Michigan or the Conference Board). The specific sample periods are discussed in Appendix 2.B and depend on the starting date of individual confidence indices. For all VARs we end the sample either in 2008:Q3 (to exclude the ZLB period on the FED Funds rate) or in 2016:Q2 (in order to also include the Great Recession and its aftermath). The results are qualitatively and quantitatively the same in both specifications.

1.4.2 Estimation

The VAR is estimated via Bayesian methods as in Uhlig (1998) and Uhlig (2005). Specifically, Uhlig's approach is followed exactly in terms of both distributional assumptions and of priors. The distributions for the VAR's coefficients and its covariance matrix are postulated to belong to the Normal-Wishart family. For estimation details, the reader is referred to either the Appendix of Uhlig (1998) or to Appendix B of Uhlig (2005). Results are based on 1,000,000 draws from the posterior distribution of the VAR's reduced-form coefficients and the covariance

matrix of its reduced-form innovations. The draws are computed exactly as in Uhlig (1998) and Uhlig (2005). The reason for using so many draws for the reduced-form VAR is that for each draw, we consider only one random rotation matrix computed by combining the methodology of Uhlig (2003) and Uhlig (2004) in order to identify permanent IS and N shocks, and the methodology of Arias, Rubio-Ramírez, et al. (2018) in order to identify the remaining shocks. We set the lag order to $p=4$. Finally, in drawing the VAR's coefficients we do not impose stationarity.

1.4.3 Identification

Our identification relies on combining two popular identification strategies in the empirical macroeconomic literature. The first strategy relies on the maximum FEV approach pioneered by Uhlig (2003) and Uhlig (2004), in which a shock is identified as the one that dominates the FEV of a particular variable at a specific horizon. The second strategy relies on the identification using robust zero and sign restrictions based on Arias, Rubio-Ramírez, et al. (2018). Note that all of our identifying restrictions have been derived from the model of ACD.⁶ The restrictions pertaining to IS and N shocks are however of more general validity and, as discussed by Fisher (2006), hold within any meaningful DSGE model. Specifically,

(i) IS shocks are postulated to be the only driver of the permanent component of the RPI.⁷ Since within the present context we are estimating the VARs in levels, we identify IS shocks as in Uhlig (2003) and Uhlig (2004) as the disturbances explaining the maximum fraction of the FEV of the RPI at a 'long' horizon, which we set to 25 years ahead.

(ii) Conditional on having identified IS shocks, we identify N disturbances based on the restriction that, among the remaining shocks, they explain the maximum fraction of the FEV of consumption at the same long horizon. Once again, the restriction holds exactly both within the ACD model, and more generally within any meaningful DSGE model.

(iii) Conditional on having identified IS and N shocks, we then identify the remaining transitory shocks by imposing a combination of zero and robust sign restrictions on impact, based on the Gibbs-sampling algorithm proposed by Arias, Rubio-Ramírez, et al. (2018). We impose zero restrictions only in population, as within ACD's model, IS shocks are the only driver of the RPI at all horizons, which implies that the impact at $t=0$ of all other shocks on the RPI has to be set equal to zero. When we work with actual data, we do not impose the zero restrictions for the RPI on impact. Instead, we leave the impacts on the RPI at $t=0$ of

⁶ACD's model features more observed variables and more structural shocks than the seven we consider here. The only motivation for uniquely focusing on seven variables and shocks is in order to avoid working with excessively large systems.

⁷Within ACD's model, IS shocks are the only disturbances impacting upon the RPI, which means that they are, in fact, the only driver of the RPI at all horizons.

all shocks other than the IS shock unconstrained.⁸ Moreover, since we will ultimately impose the restrictions on actual data, we consider it more sensible to allow for one unidentified shock. In what follows, we therefore do not impose any restrictions on the column of the structural impact matrix at $t=0$ corresponding to this disturbance. The rationale for doing so is that it is highly implausible (and in fact essentially impossible) that the disturbances we consider are the only ones impacting upon real-world economies (as opposed to the theoretical ACD model we consider). As a consequence, we leave one unidentified disturbance as a ‘catch-all’ shock, which is going to ‘hoover up’ the residual variance in the data over and above the six shocks we identify. For consistency, when we work in population, we thus identify the same six-shock specification. Note then that the only difference between the restrictions imposed in population and on real data is with regards to the zero restriction on the RPI in population, as outlined above. To summarize, beyond the permanent IS and N shocks, we consider the following transitory disturbances: a ‘pure sentiment’ shock, a monetary policy shock, a transitory N shock, a government spending shock, and a preference shock, which we however leave unrestricted.

1.4.4 Deriving the robust sign restrictions

Conceptually in line with Canova and Paustian (2011), we derive robust sign restrictions that we impose in order to identify the shocks of interest. In the context of robust sign restrictions, ‘robust’ means holding for an overwhelmingly large fraction of plausible random combinations of the model’s parameters.⁹ We consider the 90% highest posterior density intervals (HPDI) as depicted in ACD for the model’s structural parameters. Following Canova and Paustian (2011), we then take 100,000 successful draws for the parameters from a Uniform distribution defined over these intervals.¹⁰ For each draw of the parameters, we solve the model and compute impulse response functions (IRFs) to the structural shocks. The results are reported in Table 1.4, where a ‘+’ indicates a robustly positive sign, a ‘−’ a robustly negative sign and a ‘?’ indicates that the shock is left unrestricted.

⁸Note that in general, within standard DSGE models the RPI is impacted upon, at $t=0$, by all structural shocks, see e.g. the work of Justiniano et al. (2011).

⁹We use a 90% threshold. The goal of utilizing the methodology of Canova and Paustian (2011, p. 346) is to ‘(...) derive restrictions which are robust to parameter variations, independent of the specification of nuisance features, and common to the sub-models in the class to identify shocks in the data (...)’.

¹⁰The sign restrictions are the same when we use 300,000 successful draws. The Uniform distribution is chosen in order to stay agnostic about the distribution and to only make an assumption about the domain in which the parameter may lie.

Table 1.4: Robust sign restrictions on impact based on the ACD model

	ϵ_t^S	ϵ_t^M	ϵ_t^G	$\epsilon_t^{N,T}$	ϵ_t^P
Consumption	+	-	-	+	?
GDP	+	-	+	+	?
RPI	0	0	0	0	?
Hours	+	-	+	-	?
Inflation	+	-	+	-	?
Interest rate	+	+	+	-	?
Expectation of output	+	-	+	?	?

Permanent IS and N shocks are identified via Uhlig's (2003, 2004) approach. ϵ_t^S = sentiment shock; ϵ_t^M = monetary shock; ϵ_t^G = government spending shock; $\epsilon_t^{N,T}$ = transitory N shock; ϵ_t^P = preference / 'catch-all' shock. Note that the preference shock is left unrestricted, therefore we do not report the robust sign restriction but note down a '?' to indicate that the shock was left unrestricted. Also note that we leave the reaction of expectation of output to the transitory N shock unrestricted, as the sign was not robustly classified as positive or negative.

We impose the robust sign restrictions only on impact. The main reason for doing so is that, as we will show in the next sub-sections, this is already sufficient to recover all of the shocks' IRFs and fractions of FEV. Intuitively, this is connected with the fact that since we will be imposing all of the restrictions¹¹ reported in Table 1.4, we are already imposing a significant amount of information. This implies that although imposing additional restrictions at longer horizons would produce more precise results, in practice the gains would be limited. Working in population - that is, based on the theoretical MA representation of ACD's model - we now turn to the issue of whether our restrictions allow to recover the shocks' IRFs and fractions of FEV.

1.4.5 Can our restrictions recover the shocks' IRFs and fractions of FEV?

We extract from ACD's model the structural MA representation for the seven observed variables corresponding to the seven series we consider in our empirical implementation as a function of the seven structural disturbances. The seven series we consider include the RPI, GDP, consumption, hours, inflation, the monetary policy rate, and the 'individual expectation of aggregate output'. Within ACD's model, the latter plays the role of a measure of sentiment. The structural MA representation of the model can then be trivially computed based on the

¹¹With the exception of the zero restrictions on the RPI at $t=0$ for shocks other than permanent IS shocks when we work with the actual data.

model's IRFs to the seven structural shocks. Let the theoretical structural MA representation of the model's observables be

$$Y_t = A_0\epsilon_t + A_1\epsilon_{t-1} + A_2\epsilon_{t-2} + A_3\epsilon_{t-3} + \dots, \quad (1.4.2)$$

where the vectors of the observables and of the structural disturbances have been defined before, and $E[\epsilon_t\epsilon_t'] = I_N$, with I_N being the $(N \times N)$ identity matrix, so that each of the i -th columns (with $i = 1, 2, \dots, N$) of the MA matrices $A_0, A_1, A_2, A_3, \dots$ has been divided by the standard deviation of the i -th shock. Observationally equivalent reduced-form representations of (1.4.2) can be obtained by post-multiplying all of the MA matrices $A_0, A_1, A_2, A_3, \dots$ by an orthogonal rotation matrix R , yielding

$$Y_t = \tilde{A}_0\tilde{\epsilon}_t + \tilde{A}_1\tilde{\epsilon}_{t-1} + \tilde{A}_2\tilde{\epsilon}_{t-2} + \tilde{A}_3\tilde{\epsilon}_{t-3} + \dots, \quad (1.4.3)$$

where $\tilde{A}_j = A_jR$ and $\tilde{\epsilon}_{t-j} = R'\epsilon_{t-j}$, $j = 0, 1, 2, 3, \dots$. We can randomly generate different rotation matrices,¹² thus producing different observationally equivalent vector moving average (VMA) representations. The question we wish to address is whether imposing the previously discussed identifying restrictions on different reduced-form VMAs allows us to recover the true IRFs and fractions of FEV.

In order to qualify our results, we perform three exercises distinct only in their parameterization of the ACD model. As a baseline, we use ACD's median estimates of the 90% HPDI interval for their confidence augmented sticky price model. In the first exercise, ('Experiment 1'), we then perform the estimation using as parameter values said median estimates. These parameters include relevant levers such as the persistence of confidence shock (ρ_ξ), the standard deviation of the confidence shock (σ_ξ), and the standard deviation of the IS shock (σ_{ip}).¹³ In the second exercise, ('Experiment 2'), we also use the baseline parameterization as well as the median estimates for ρ_ξ and σ_{ip} . For σ_ξ however, we use 80% of its median estimate.¹⁴ Finally, in the third exercise ('Experiment 3'), we perform the estimation using again the baseline parameterization, modulo the lower bound of the 90% HPDI interval for ρ_ξ and σ_ξ , as well as the upper bound for σ_{ip} .¹⁵

The motivation behind our experiments is as follows. In Experiment 1, we want to see how our identification strategy performs when we take ACD's model at face value. Given this, in Experiment 2, we want to evaluate the performance when the confidence shock is forced to be 'less important', that is, by decreasing its standard deviation. Finally, in Experiment 3, we

¹²We generate the $(N \times N)$ random rotation matrix as follows. We start by taking an $(N \times N)$ draw K from a $\sim N(0, 1)$ distribution. Then, we take the QR decomposition of K , and we set the rotation matrix to Q' .

¹³Hence, $\sigma_{ip} = 0.610$, $\rho_\xi = 0.833$, and $\sigma_\xi = 0.613$.

¹⁴Hence, $\sigma_{ip} = 0.610$, $\rho_\xi = 0.833$, and $\sigma_\xi = 0.613 * 0.8 = 0.490$.

¹⁵Hence, $\sigma_{ip} = 1.306$, $\rho_\xi = 0.717$, and $\sigma_\xi = 0.348$.

want to evaluate our identification when (i) the confidence shock is less important and (ii) the IS shock is more important, since - as we will see in the empirical application in Section 1.5 - there, IS and N shocks in fact explain a dominant fraction of the FEV of all variables.

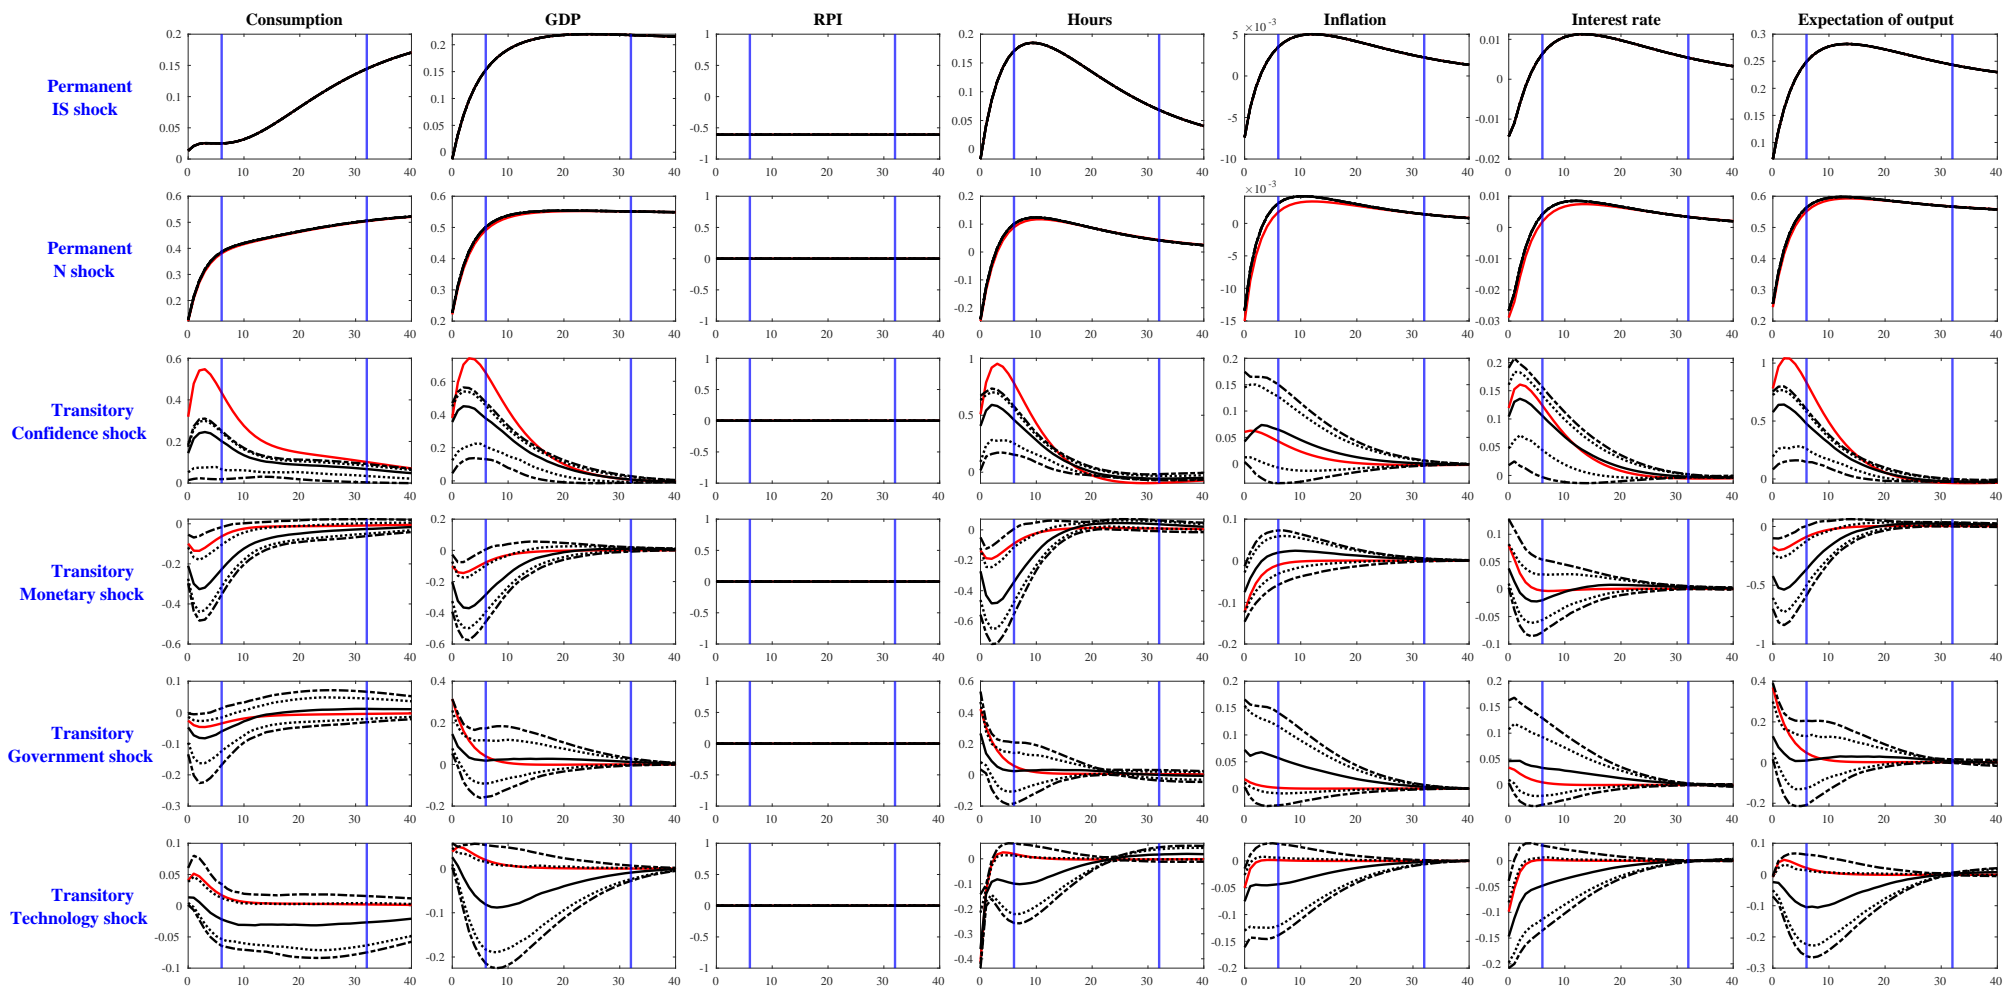
Figures 1.3 - 1.8 report the results of our experiments, for the IRFs, and the fractions of FEV, respectively. The vertical blue lines depict the 6th and the 32nd quarter ahead to denote the band generally associated with business-cycle fluctuations. In all figures, the objects pertaining to IS and N shocks (IRFs, or fractions of FEV) are exactly recovered, reflecting the fact that for either shock, the maximum FEV approach produces a single matrix for each individual reduced form VMA representation. As for the other shocks, the presence of rotation uncertainty originating from the algorithm of Arias, Rubio-Ramírez, et al. (2018) (which we use to jointly impose the zero and sign restrictions) implies that we will have distributions for the IRFs and fractions of FEV, rather than point estimates.

We find that in all experiments, the shocks pertaining to the transitory monetary, government, and N shocks are recovered with precision as in the majority of cases, the 16-84 percentiles contain the true IRFs and fractions of FEV at all horizons and the 5-95 percentiles contain the true objects in all cases. Regarding the identification of the confidence shock, the following pattern emerges. In Experiment 1, we can see that the confidence shock is mainly underestimated for consumption, with the true IRFs and FEV outside the 5-95 percentile at all horizons. For GDP and hours worked, the true IRFs and FEVs are for the first few quarters within the 16-84 percentile confidence bands, but outside thereafter. Only after about 15 quarters do the true IRFs overlap again with the 5-95 percentiles. The results of Experiment 2 resemble those of Experiment 1, with a smaller difference between the true and estimated IRFs and FEVs. In Experiment 3, the confidence shock is recovered for GDP and hours worked with great precision. For consumption, the true IRFs lie within the 16-84 percentiles and match the median at around 12 quarters, while the FEV of the confidence shock for consumption is still underestimated.

The above experiments drive home the following three points: First, our identification approach always identifies IS and N shocks either exactly or with great precision. Second, when the IS shock explains only very little of the FEVs of the variables considered and the confidence shock explains a large portion of the FEVs of all variables instead, (i.e., as in the model by ACD), then our methodology underestimates the IRFs and FEVs of the confidence shock, but still recovers all other shocks precisely. Third, when the importance of IS and N shocks is closer to what we observe in the empirical exercise in the next section, then the identification of all shocks - including the confidence shock - works very well, with the only slight exception being consumption.


Experiment 1

Figure 1.3: Experiment 1, IRF



Recovering the IRFs of the ACD model in population (red: true IRFs; black: median, 16-84, and 5-95 percentiles of the distribution of estimated IRFs).

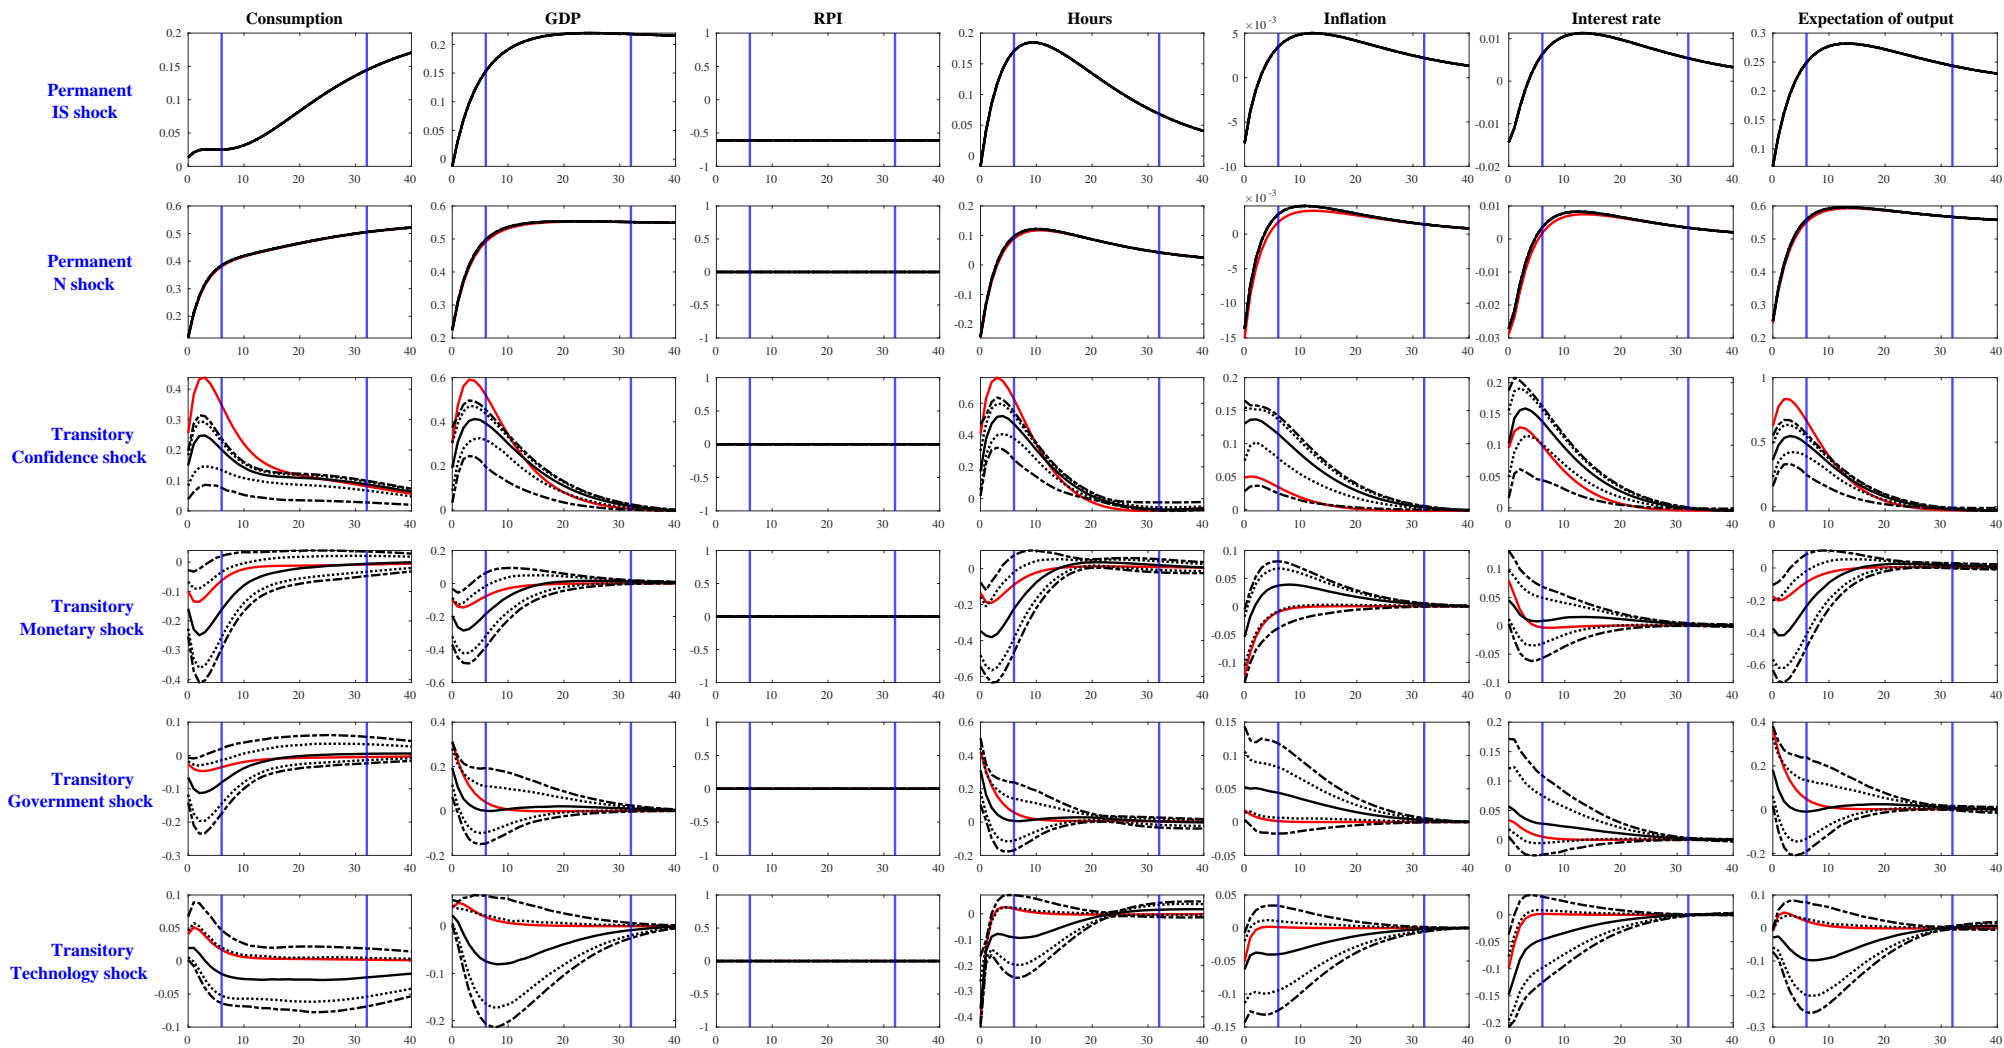
Figure 1.4: Experiment 1, FEV



Recovering the FEVs of the ACD model in population (red: true FEVs; black: median, 16-84, and 5-95 percentiles of the distribution of estimated FEVs).


Experiment 2

Figure 1.5: Experiment 2, IRF



Recovering the IRFs of the ACD model in population (red: true IRFs; black: median, 16-84, and 5-95 percentiles of the distribution of estimated IRFs).

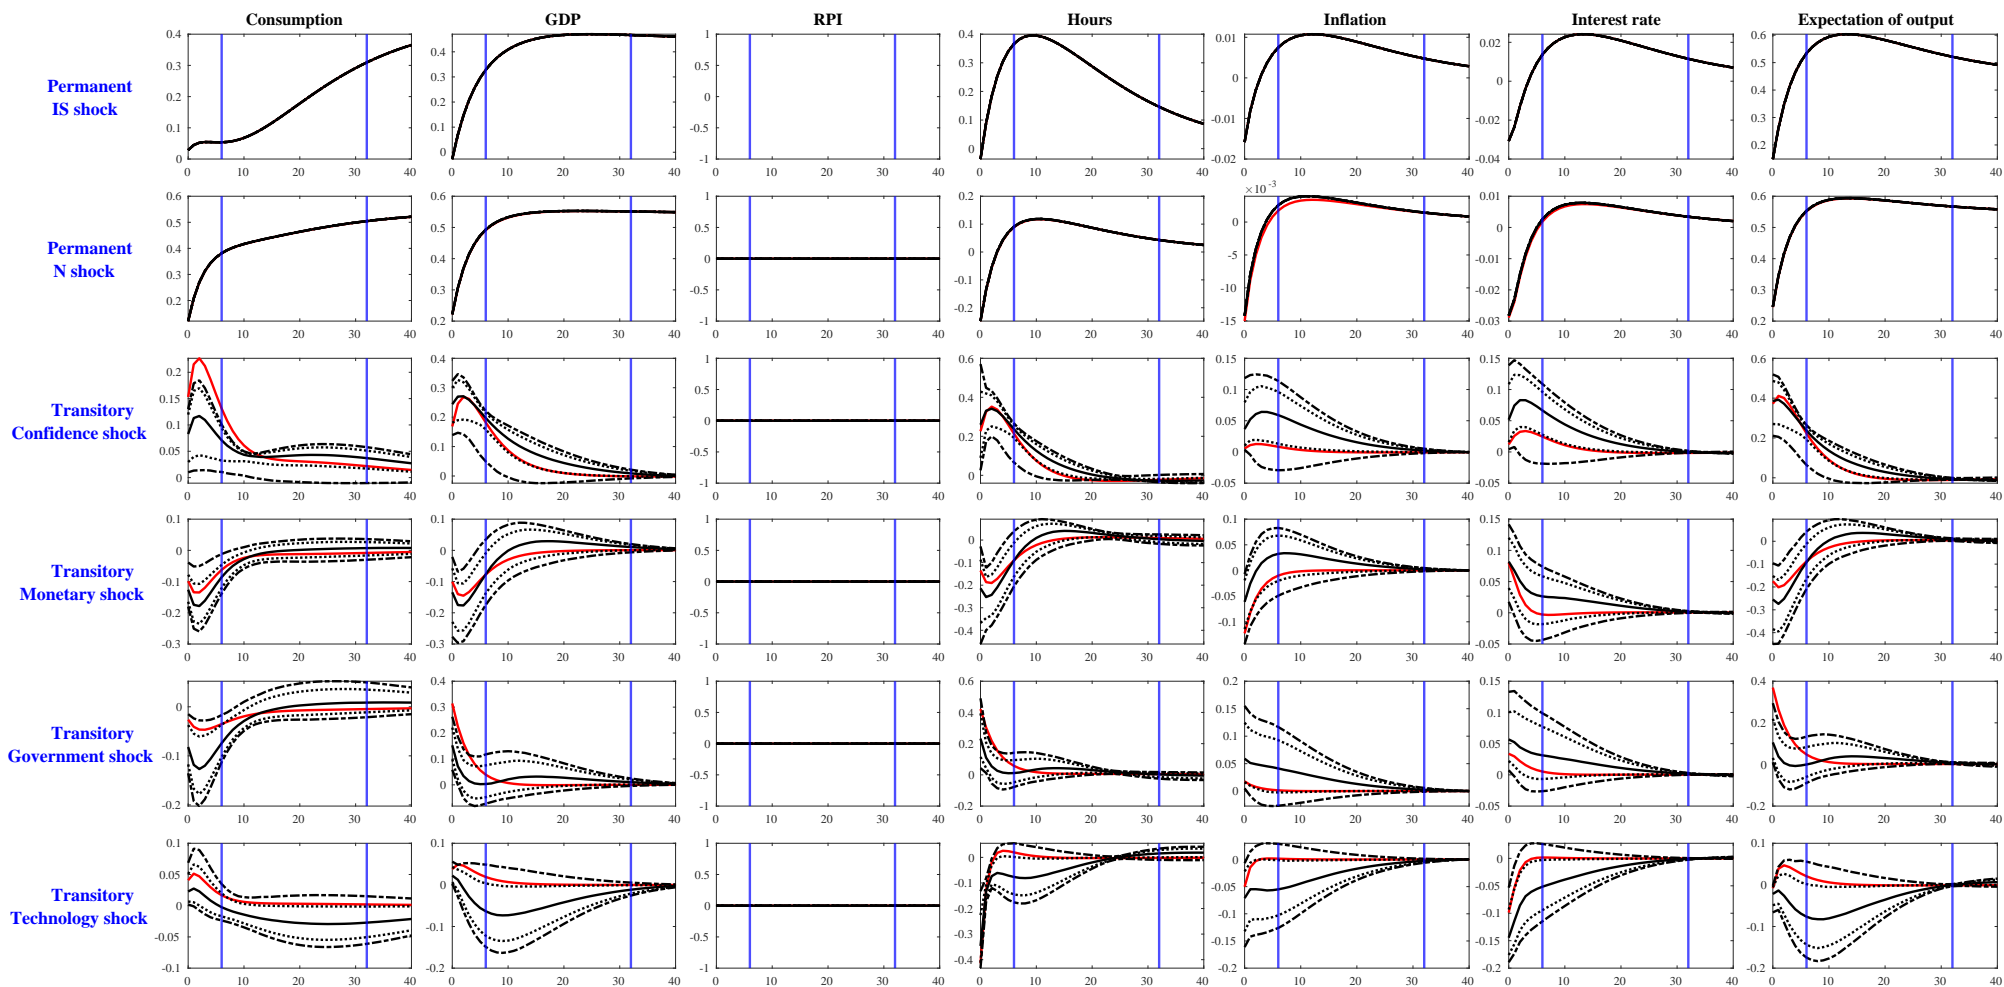
Figure 1.6: Experiment 2, FEV



Recovering the FEVs of the ACD model in population (red: true FEVs; black: median, 16-84, and 5-95 percentiles of the distribution of estimated FEVs).


Experiment 3

Figure 1.7: Experiment 3, IRF



Recovering the IRFs of the ACD model in population (red: true IRFs; black: median, 16-84, and 5-95 percentiles of the distribution of estimated IRFs).

Figure 1.8: Experiment 3, FEV



Recovering the FEVs of the ACD model in population (red: true FEVs; black: median, 16-84, and 5-95 percentiles of the distribution of estimated FEVs).

1.5 Evidence

Having been reassured that our identifying restrictions allow us to recover - exactly, for permanent IS and N shocks, and with reasonable precision, for all other disturbances - the shocks' true IRFs and fractions of FEV, we now turn to the actual data. As outlined previously, the only difference between the restrictions we just imposed upon the theoretical reduced-form MA representation of ACD's model, and those we will impose upon the actual data pertains to the impacts on the RPI at $t=0$ of all shocks other than permanent IS and N disturbances. Whereas when working with the theoretical MA representation of ACD's model these impacts were restricted to zero, when we work with the actual data we will leave them unrestricted.

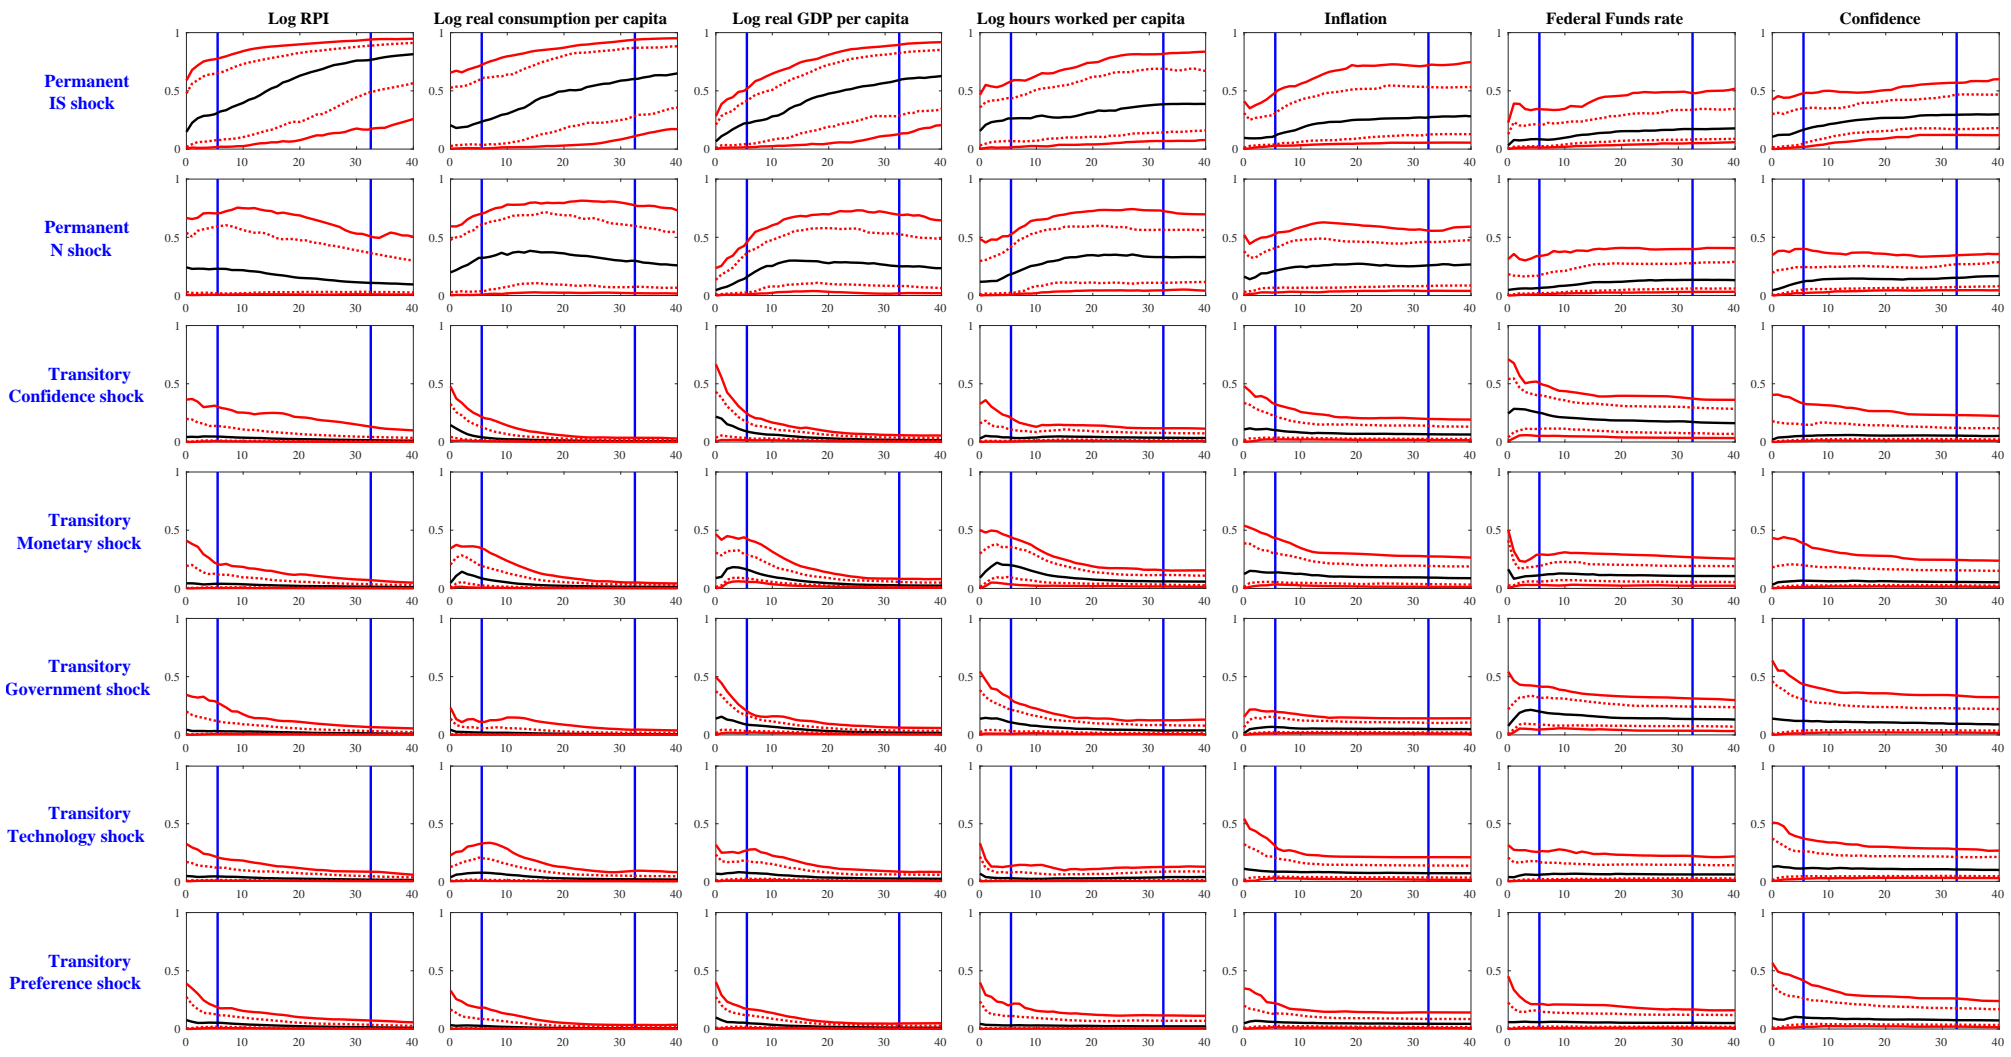
Figures 1.9 and 1.10 show the fractions of FEV at horizons up to 10 years ahead explained by the seven shocks, based on systems including the Conference Board overall index and the University of Michigan's index of consumer sentiment pertaining to the question: 'Now turning to business conditions in the country as a whole - do you think that during the next 12 months we'll have good times financially, or bad times, or what?'¹⁶ Figure 1.11 reports the same evidence from a different perspective. It shows the estimated fractions of FEV jointly explained by permanent IS and N shocks, and the residual fractions of FEV. The residual fractions of FEV represents the upper bound to the fractions of FEV pure sentiment shocks could explain in the circumstance in which they were the only other shock driving the economy, beyond permanent IS and N shocks. The set of results reported in Figures 1.9 - 1.11 are representative of the overall set of results based on any of the confidence indices we consider. We only report results based on these two indices for reasons of space. The main substantive findings emerging from the three figures are that first, the identified sentiment shocks explain small-to-negligible fractions of the FEV of all variables, including the confidence indices themselves, thus suggesting that measures of consumer confidence are driven, to a dominant extent, by disturbances other than pure sentiment. Second, permanent IS and N shocks jointly explain very large-to-dominant fractions of the FEV of all series at the business-cycle band of 6-32 quarters, which in Figures 1.9 - 1.11 are marked by the vertical blue lines. As discussed, (i) the presence of these two disturbances is essentially unquestioned in the macroeconomic profession; (ii) the way we identify them, via Uhlig (2003) and Uhlig (2004)'s approach, is standard; and (iii) our restrictions allow us to exactly recover the two shocks in population. The fact that these two shocks then jointly explain large-to-dominant portions of the FEV of all variables between 6-32 quarters puts a robust upper bound on the role pure sentiment shocks might play. Put differently, since permanent IS and N shocks are unquestionably there, and they play a large role in driving business-cycle fluctuations, even in the circumstance in which sentiment shocks explained all of the residual FEV of macroeconomic variables not explained

¹⁶Specifically, as in Barsky and Sims (2012), the index we are using has been computed as the difference between the 'Good times' and 'Bad times' percentages of answers.

by IS and N shocks, this would not amount to much.

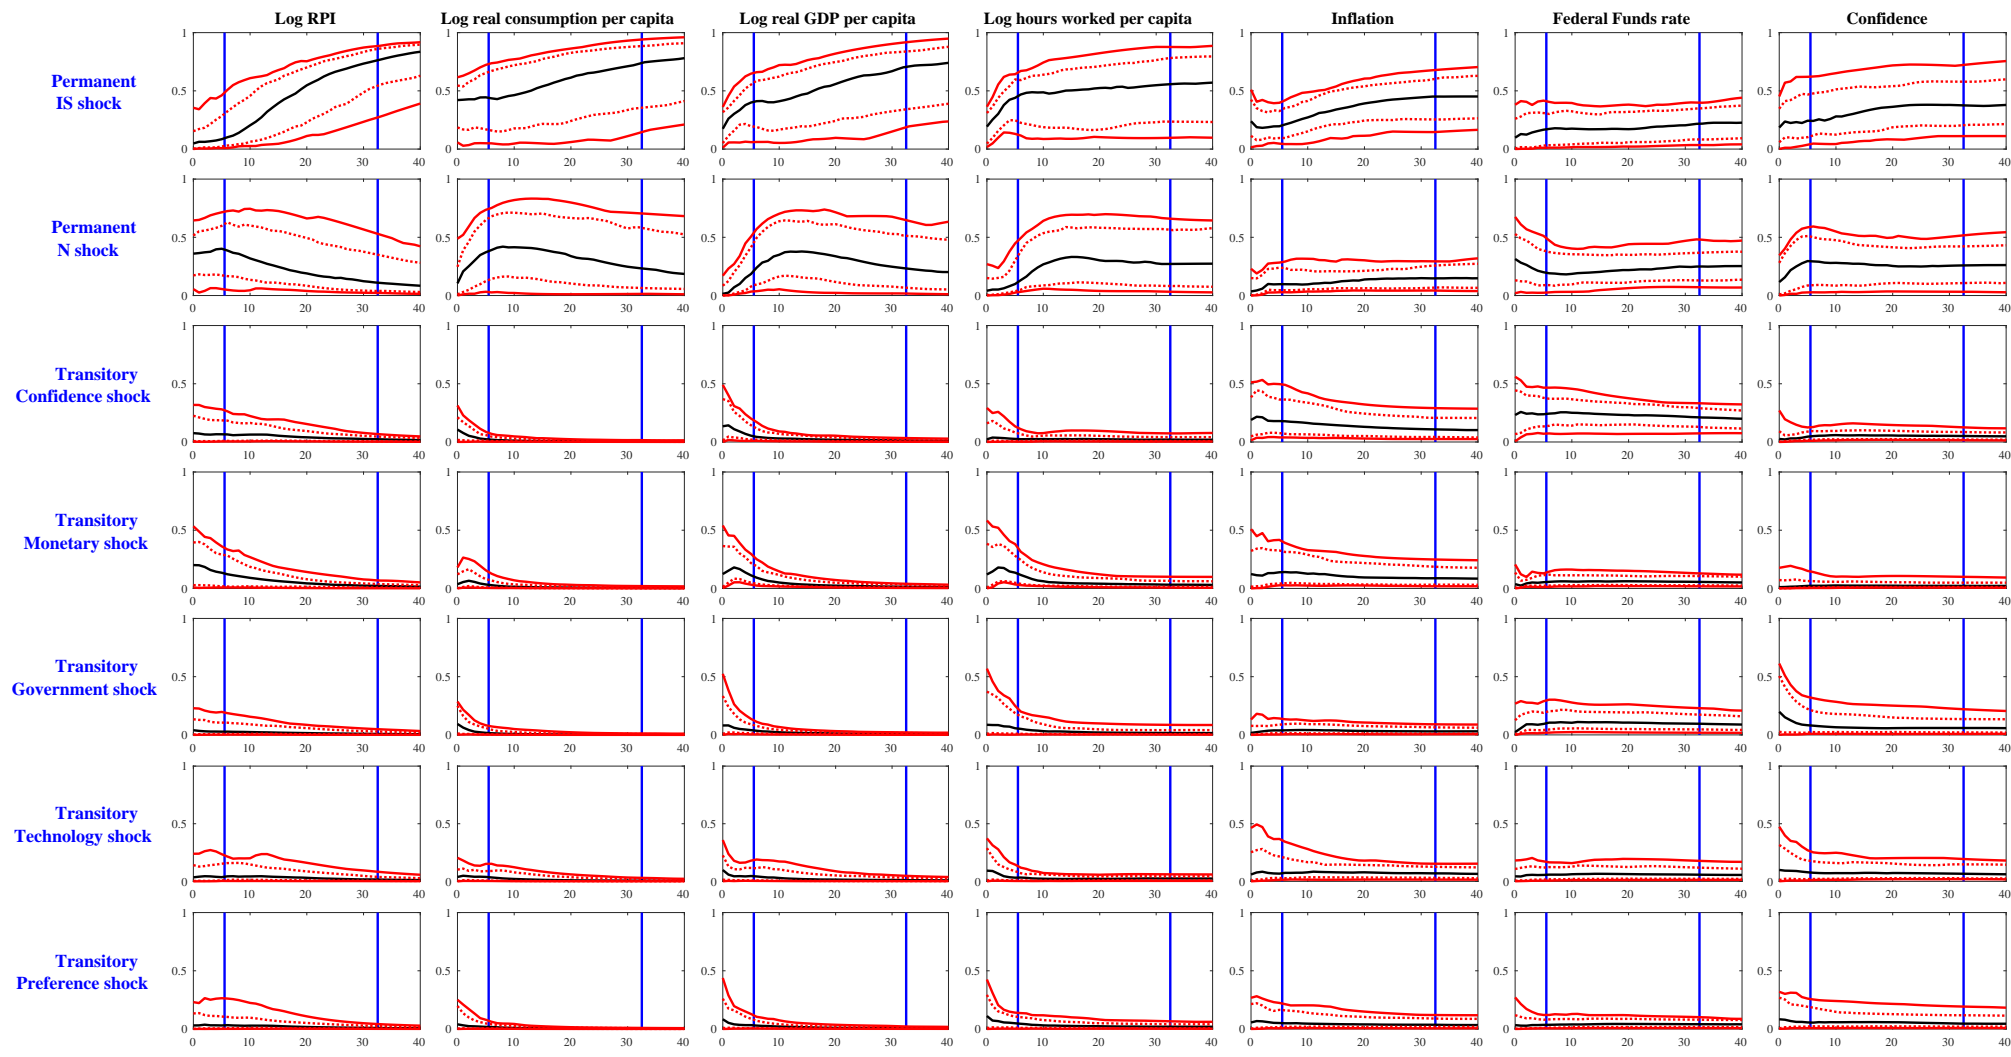
Our own conclusion is therefore that autonomous fluctuations in sentiment - even if they truly are there - play a minor-to-negligible role in macroeconomic fluctuations. This is in line with, e.g., Barsky and Sims (2012)'s conclusion that '[a]nimal spirits shocks account for negligible shares of the forecast error variances of consumption and output at all frequencies', and with Fève and Guay (2019)'s finding that the shock explaining the largest share of the residual FEV of consumer confidence indices - once having preliminarily identified permanent technology shocks - explains very little of macroeconomic fluctuations.

Figure 1.9: Real data: Conference Board Overall Index



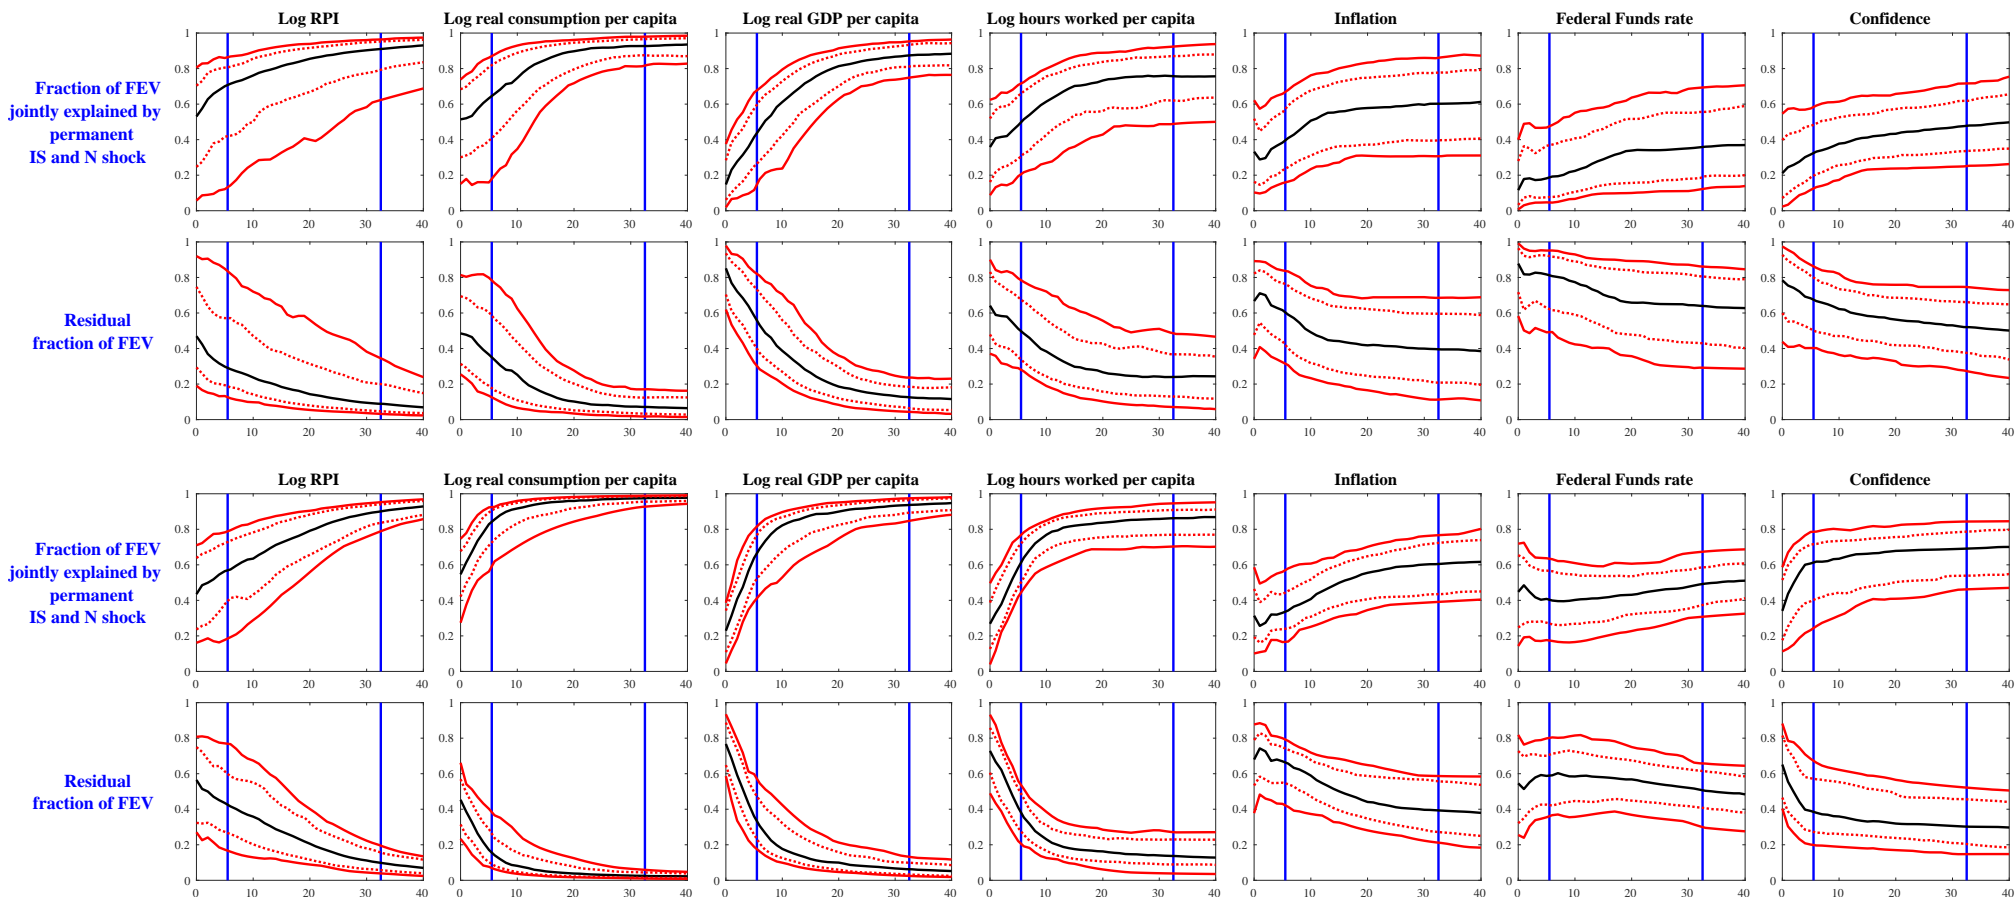
Estimated fractions of FEV (median, 16-84, and 5-95 percentiles of the posterior distribution).

Figure 1.10: Real data: Michigan Index 10



Estimated fractions of FEV (median, 16-84, and 5-95 percentiles of the posterior distribution).

Figure 1.11: Estimated fractions of FEV jointly explained by permanent IS and N shocks, and residual fraction of FEV



Estimated fractions of FEV (median, 16-84, and 5-95 percentiles of the posterior distribution).

1.6 Conclusion

In this paper we have made two contributions to the literature exploring the role of sentiment in macroeconomic fluctuations. First, working with the theoretical MA representations of standard DSGE models, we have shown that several SVAR-based approaches to the identification of sentiment shocks are unreliable, as they identify such disturbances even when the model does not feature them. The approach proposed by Beaudry, Nam, et al. (2011), for example, identifies sentiment shocks within SW's model. The problem is that the restrictions which are typically imposed are so weak and generic that they will always be satisfied with non-negligible probability by random rotations of the model's structural disturbances, irrespective of the fact that they do or do not include a pure sentiment shock. Second, we have derived robust restrictions for the identification of sentiment shocks based on the model of ACD, and working with the theoretical MA representation of the model, we have evaluated their performance in recovering the shocks' true IRFs and fractions of FEV. We find that our restrictions always allow us to recover permanent IS and N shocks exactly. When we impose our restrictions upon US macroeconomic data within a SVAR framework, we find that permanent IS and N shocks already jointly explain large-to-dominant portions of the FEV of all real macroeconomic variables between 6-32 quarters and detect only a minor-to-negligible role for sentiment shocks in business-cycle fluctuations.

1.A Appendix I

1.A.1 Macroeconomic data

The following series are all available at the quarterly frequency:

- John Fernald’s ‘purified TFP’ series is available from the San Francisco Fed’s website.
- Real output per hour of all persons in the non-farm business sector (OPHNFB) is from the U.S. Bureau of Labor Statistics.
- A seasonally adjusted series for real GDP (GDPC96) is from the U.S. Department of Commerce: Bureau of Economic Analysis.
- Inflation has been computed as the log-difference of the GDP deflator (GDPCTPI) taken from the St. Louis Fed’s website.
- Hours worked by all persons in the non-farm business sector (HOANBS) is from the U.S. Department of Labor, Bureau of Labor Statistics.
- The seasonally adjusted series for real chain-weighted investment, consumption of non-durables and services, and their deflators (which we use to compute the chain-weighted relative price of investment) have been computed based on the data found in Tables 1.1.6, 1.1.6B, 1.1.6C, and 1.1.6D of the National Income and Product Accounts. Real consumption and its deflator pertain to non-durables and services. Real investment and its deflator have been computed by chain-weighting the relevant series pertaining to durable goods; private investment in structures, equipment, and residential investment; Federal national defense and non-defense gross investment; and State and local gross investment.

The remaining variables are available at a monthly frequency and have been converted to the quarterly frequency by taking averages within the quarter.

- The Federal funds rate and the 1-, 3-, 5-, and 10-year government bond yields are from the St. Louis Fed’s website. They are quoted at a non-annualized rate to make their scale exactly comparable to that of inflation.¹⁷
- Civilian non-institutional population (CNP16OV) is from the U.S. Department of Labor, Bureau of Labor Statistics.
- The BAA-AAA spread is calculated from Moody’s Seasoned Aaa Corporate Bond Yield (AAA) and Moody’s Seasoned Baa Corporate Bond Yield (BAA) from the Board of Governors of the Federal Reserve System.

¹⁷If we define an interest rate series as R_t —with its scale such that, e.g., a ten percent rate is represented as 10.0—the rescaled series is computed as $r_t = (1 + R_t/100)^{1/4} - 1$.

1.A.2 Confidence indices

Here follows a description of the consumer and CEO indices from the University of Michigan and the Conference Board.

Indices from the University of Michigan

From the University of Michigan website, we took the following consumer confidence indices. In what follows, for ease of reference, we refer to them as ‘Index 1’, ‘Index 2’, etc.

- *Index 1*: ‘Overall Index’ (from Table 1: The Index of Consumer Sentiment)
- *Index 2*: Index component ‘Current Index’ (from Table 5: The Components of the Index of Consumer Sentiment)
- *Index 3*: Index component ‘Expected Index’ (from Table 5: The Components of the Index of Consumer Sentiment)
- *Index 4*: Index component ‘Personal Finances, Expected’ (from Table 5: The Components of the Index of Consumer Sentiment)
- *Index 5*: Index component ‘Business Conditions 12 Months Ahead’ (from Table 5: The Components of the Index of Consumer Sentiment)
- *Index 6*: Index component ‘Business Conditions 5 Years Ahead’(from Table 5: The Components of the Index of Consumer Sentiment)
- *Index 7*: The question was: ‘*Now looking ahead – do you think that a year from now you (and your family living there) will be better off financially, or worse off, or just about the same as now?*’ Computed as better off minus worse off, from Table 8: Expected Change in Financial Situation in a Year.
- *Index 8*: The question was: ‘*During the last few months, have you heard of any favorable or unfavorable changes in business conditions? What did you hear?*’ Computed as favorable news minus unfavorable news, from Table 23: News Heard of Recent Changes in Business Conditions.
- *Index 9*: The question was: ‘*And how about a year from now, do you expect that in the country as a whole business conditions will be better, or worse than they are at present, or just about the same?*’ Computed as better minus worse, from Table 26: Expected Change in Business Conditions in a Year.

- *Index 10*: The question was: ‘*Now turning to business conditions in the country as a whole – do you think that during the next 12 months we’ll have good times financially, or bad times or what?*’ Computed as good times minus bad times, from Table 28: Business Conditions Expected During the Next Year.
- *Index 11*: The question was: ‘*Looking ahead, which would you say is more likely – that in the country as a whole, we’ll have continuous good times during the next 5 years or so, or that we will have periods of widespread unemployment or depression, or what?*’ Computed as good times minus bad times, from Table 29: Business Conditions Expected During the Next 5 Years.

Indices 1 to 3 are available since 1960Q1. Indices 4 to 7, 10, and 11 are available since 1959Q4. Index 8 is available since 1965Q1. Index 9 is available since 1965Q3.

Confidence indices from the Conference Board

Allen Li of the Conference Board has kindly provided the following confidence indices. For ease of reference, we refer to them as ‘Index 1’, ‘Index 2’, etc.

- *Index 1*: Consumer Confidence Index[®] : Overall index
- *Index 2*: Consumer Confidence Index[®] : Present Situation
- *Index 3*: Consumer Confidence Index[®] : Expectations
- *Index 4*: Measure of CEO Confidence[™]: Overall index
- *Index 5*: Measure of CEO Confidence[™]: Current Economic Conditions vs. 6 Months Ago
- *Index 6*: Measure of CEO Confidence[™]: Expectations for Economy, 6 Months Ahead
- *Index 7*: Measure of CEO Confidence[™]: Expectations for Own Industry 6 Months Ahead

Indices 1 to 3 are available since 1978Q1. Indices 4 to 7 are available since 1976Q2.

Chapter 2

Business-Cycle Anatomy around the World

Abstract

Are business-cycles all alike and driven by the same type of main force in the cross-section of advanced economies? We apply the strategy of Angeletos et al. (2020) for dissecting macroeconomic time series to analyze the main business-cycle (MBC) driver in Australia, Canada, France, Italy, the United Kingdom (UK), and the United States (US). Our results support the existence of an MBC, a single dominant force in each country that explains the bulk of macroeconomic fluctuations, triggers strong comovement in the main macroeconomic variables, and a countercyclical reaction of net exports as imports increase by more than exports. The identified MBCs play a significant role in explaining international output comovement and are disconnected from inflation, the terms of trade, and the supply forces that drive economic activity in the long-run. Our findings support Lucas (1977) and indicate that business-cycles across the studied advanced economies are all alike and driven by the same type of non-inflationary, domestic aggregate demand force.

Acknowledgment: I am sincerely grateful to Fabrice Collard and Harris Dellas for their continuous support and invaluable guidance throughout the duration of this project. Moreover, I thank Luca Benati, Pierpaolo Benigno, Marco D'Amico, Lorenz Driussi, Annina Heini, Sylvia Kaufmann, Armando Näf, Dirk Niepelt, Aubrey Poon, Fabienne Schneider, and Larissa Schwaller for helpful comments and suggestions. Furthermore, I am grateful for helpful comments from seminar and conference participants in the Macro Group Meeting and Macro Ph.D. Seminar at the Department of Economics at the University of Bern.

2.1 Introduction

Lucas (1977) famously observed that ‘business-cycles are all alike’ and that this ‘suggests the possibility of a unified explanation of business-cycles’. While national, unconditional variance-covariance matrices of the cyclical components of macroeconomic aggregates are similar across advanced economies, similar unconditional moments alone are not enough to assume that business-cycles are all alike. Such conformity could result from different shocks with distinct propagation mechanisms driving macroeconomic fluctuations in different countries. This would imply that business-cycles are not fully alike along the international dimension and, as a result, cannot be studied using a unifying business-cycle theory. Recently, Angeletos et al. (2020) (henceforth ACD) introduced the concept of a ‘main business-cycle’ (MBC) shock to refer to the force that accounts for the bulk of business-cycle fluctuations in the main macroeconomic variables in the United States (US).

Our key objective in this paper is to investigate whether similar main drivers of business-cycles exist in the cross-section of advanced countries to assert that business-cycles are indeed all alike and driven by the same type of shock. We identify the main drivers of business-cycle fluctuations in six advanced economies, including Australia (AU), Canada (CA), France (FR), Italy (IT), the United Kingdom (UK), and the United States, to establish the existence of a common empirical template. Additionally, we examine the comovement of the identified forces across countries and study the relationship between the main domestic business-cycle driver and variables related to international trade. Our identification strategy follows ACD, which relies on a ‘max-share identification’ approach. With ACD’s approach, a shock is identified to explain the maximum volatility of a particular variable over a specific frequency band, e.g., business-cycle or long-run frequencies.

Summarizing the results, we find support for the existence of a single dominant force - the MBC shock - in each of the countries above that explains the bulk of macroeconomic fluctuations, triggers strong comovement in the main macroeconomic variables, and is disconnected from inflation and the supply forces that drive economic activity in the long-run. In terms of international synchronization, the domestic MBC shocks explain a considerable fraction of international output comovement. Concerning the link between the main driver of fluctuations in aggregate activity and international trade, we find a strong connection between the MBC shock and imports, as well as a strong connection of the main driver of imports and the key macroeconomic variables. The main driver of exports also shares similar dynamics as the MBC shock, but has a more narrow footprint focused on output. Similar to the disconnect between the MBC shock and inflation, we observe a disconnect between the MBC shock and the terms of trade and between the main driver of the terms of trade and the key macroeconomic variables. The main driver of fluctuations in aggregate activity and the main driver of the terms of trade

are disconnected and representatives of distinct forces.

Our paper contributes to the literature in three ways. First, it extends ACD’s closed economy evidence for the US to open economies and provides a common empirical template of a main business-cycle driver across countries, echoing Lucas (1977) in that business-cycles are truly all alike and driven by the same type of shock. Second, it implies that demand-side shocks are important drivers of international business-cycle comovement, in line with recent research in the international business-cycle literature (Levchenko and Pandalai-Nayar (2020) and Huo et al. (2020)). Third, it provides food for thought for parsimony-inclined theoretical models by highlighting the need for different shocks that account for business-cycle fluctuations in the macroeconomic aggregates and the terms of trade, in line with recent evidence that shows a similar disconnect between the real exchange rate and the main driver of domestic macroeconomic activity (Miyamoto et al. (2022)).

The remainder of the paper is organized as follows: Section 2.2 summarizes the key literature. Section 2.3 describes the data and the empirical strategy. Section 2.4 discusses the results of the core set of domestic macroeconomic variables and the degree of international synchronization. Section 2.5 discusses the results regarding the inclusion of international trade variables. Section 2.6 concludes.

2.2 Literature

Our paper relates to three strands of the literature. First, it contributes to the vast macroeconomic literature investigating the drivers of business-cycle fluctuations. Starting with Sims (1980), a common approach in this literature is to estimate vector autoregressions (VARs) or vector error correction models (VECMs), make assumptions to identify a structural shock, and to characterize its effect on macroeconomic aggregates using impulse response functions (IRFs) and forecast error variance decompositions (FEVDs). An IRF summarizes the time path of the dynamic causal effect of a structural shock on a macroeconomic variable. The FEVD quantifies the relative importance of a structural shock in explaining the variation of a macroeconomic variable. Structural shocks are defined as ‘unexpected exogenous disturbances to structural economic relationships’ (Stock and Watson (2016)). Emphasized proponents include technology shocks, monetary policy shocks, fiscal shocks, oil shocks, and, more recently, uncertainty shocks, financial shocks, and confidence shocks. An overview of common restrictions, identification schemes, and identified shocks is given in Ramey (2016). In this paper, we do not identify a particular structural shock but aim to identify the *main* driver of a country’s business-cycle to establish a common empirical template across countries. We apply the identification strategy of ACD, which builds on the max-share identification approach pioneered by Uhlig (2003). In this approach, a shock is identified as the one that explains the dominant fraction

of the forecast error variance (FEV) of a particular variable at a given horizon. Relative to the original max-share identification, ACD use spectral decomposition to target the variance of a particular variable in a specific frequency band, allowing to isolate fluctuations pertaining to business-cycle or long-run phenomena. Unlike structural identification, ACD's approach does not impose restrictions motivated by economic theory (e.g., sign restrictions), nor does it recover a structural shock. Instead, it identifies a shock as the one that contributes the most to the volatility of a particular variable over a specific frequency band.

Second, this paper relates to the literature on dynamic factor models (DFMs), which is overviewed by Stock and Watson (2016). DFMs summarize the dynamic behavior of a possibly large vector of observed time series into a few unobserved common factors that evolve over time. ACD's identification strategy also aims to find the latent factor that drives the comovement of key macroeconomic variables. Unlike DFMs, ACD's procedure estimates a VAR or VECM to systematically vary the targeted variable and frequency band to build up a collection of shocks, each of which explains the maximum volatility of a particular variable over a specific frequency band. This multidimensional anatomy of the data provides a rich set of cross-variable, static, and dynamic properties in terms of IRFs and FEVDs, from which to derive economic meaning and inform theory about the drivers of business-cycle or long-run fluctuations.

Third, our paper adds to the literature on international business-cycles (IBCs) going back to Backus, Kehoe, and Kydland (1992), Backus and Kehoe (1992), and many others. Akin to the IBC literature, we investigate the comovement between main business-cycle drivers across countries and analyze the relationship between domestic business-cycles and fluctuations in international trade. We extend ACD's closed economy analysis to a panel of small open economies and international trade variables. Within the recent IBC literature, our paper is related to the work that emphasizes the importance of non-technology shocks to explain international business-cycle comovement instead of the common approach of working with fluctuations driven by total factor productivity (TFP). Levchenko and Pandalai-Nayar (2020) study the international propagation of business-cycles between the US and Canada. They identify a non-technology business-cycle shock for the US. Similar to our findings regarding the main drivers of business-cycles, they find that their identified transitory demand shock generates positive comovement in main macroeconomic aggregates and accounts for the majority of US business-cycle fluctuations as well as a large share of business-cycle comovement between the US and Canada. Similarly, Huo et al. (2020) estimate annual utilization-adjusted TFP series for a large group of countries. They find that utilization-adjusted TFP is uncorrelated across countries. They show that non-technology shocks that affect factor utilization are more important drivers of international business-cycles than technology shocks. Finally, Miyamoto et al. (2022), the paper most similar to ours, focuses on the connection between the real exchange rate and key macroeconomic variables. They also use ACD's identification strategy, establish the

existence of a main business-cycle driver in each of the G7 countries (US, Canada, Japan, UK, Germany, France, Italy), and find that these main drivers are highly correlated across countries. In addition to Miyamoto et al. (2022), we establish the connection of the main driver of economic activity to capacity utilization, its disconnect to technology, and the existence of a main long-run shock disconnected from the main business-cycle driver. Unlike them, we incorporate the unemployment rate and labor productivity in our baseline datasets. We include unemployment because it is considered one of the most important and recognized indicators of the state of the economy. The inclusion of labor productivity allows us to elaborate on the relationship between the main driver of economic activity and technology. Finally, we also estimate VECMs in addition to VAR models to account for cointegration relationships among variables explicitly. Similar to our finding of the disconnect between the main business-cycle driver and the terms of trade, Miyamoto et al. (2022) find a disconnect to the real exchange rate as well as a weak connection between the real exchange rate and real exports and imports. They conclude that separate shocks are needed to explain both business-cycles and real exchange rates in open economy macro models.

2.3 Empirical Strategy

In this section, we present the data and the empirical strategy. Our empirical strategy closely follows that of ACD. Relative to the traditional max-share identification on which ACD build, two distinct features stand out. First, ACD utilize spectral decomposition to isolate the variance of a specific variable within a specific frequency band, which allows for an examination of the relative contributions of different components of the data in isolation. We consider two frequency bands, one associated with business-cycle (6-32 quarters) and one associated with long-run (80- ∞ quarters) frequencies.¹ Second, ACD's strategy consists of taking multiple cuts of the data by varying the targeted variable and frequency band. These multidimensional cuts of the data help characterize the properties of the main business-cycle driver we aim to identify.

2.3.1 Data

We use quarterly observations on log real *per capita* levels of GDP (Y_t), investment (I_t), and consumption (C_t); log hours worked per person (h_t), the unemployment rate (u_t), log labor productivity (LP_t), the inflation rate (π_t), and the nominal interest rate (R_t). These eight variables constitute the baseline domestic dataset for our sample of advanced economies comprising Australia, Canada, France, Italy, the United Kingdom, and the United States. In our extended

¹Another possible choice not investigated are medium-run frequencies as studied in Comin and Gertler (2006).

dataset, we also include log real *per capita* levels of exports (EXP_t) and imports (IMP_t); and the log terms of trade (TOT_t). Including these international trade variables allows us to characterize the link between the main driver of fluctuations in aggregate activity and trade. We use data from the Organisation for Economic Co-operation and Development Economic Outlook (OECD EO) database. The inflation rate is measured as the rate of change in the GDP deflator. Investment is defined as gross capital formation (GCF) and includes changes in inventories.² Consumption is measured as the private final consumption expenditure. Due to data availability, the measurements of investment and consumption differ from their counterparts in ACD, where investment includes consumer expenditure on durables, and consumption consists of expenditure on nondurables and services. Our results for the US relative to ACD are not significantly impacted by these differences. The terms of trade are defined as the ratio between export and import prices. Since quarterly, long-horizon utilization-adjusted TFP series are currently only available for the US, TFP is not included in any of the specifications. The individual samples in the OECD EO dataset start in 1978:Q3 for AU, 1981:Q1 for CA, 1960:Q2 for FR, 1971:Q1 for IT, 1978:Q1 for the UK, and 1960:Q2 for the US. Data availability of the above macroeconomic time series in the OECD EO dataset dictates the different sample lengths. We end all samples in 2019:Q1 and thereby exclude the Covid-19 pandemic. Table 2.17 in the Appendix displays the trade openness index for each country. Due to their trade openness, AU, CA, FR, IT, and the UK are generally referred to as (small) open economies (SOE), while the US is considered a closed economy. We include the US in our sample of predominantly open economies to compare results across open and closed advanced economies and to compare the results for the US obtained in this paper to those obtained in ACD. The raw data, the construction of the series, and data sources are described in detail in Appendix 2.B.

For robustness checks, we also include the following data. First, we consider quarterly series for hours, employment, and population from the updated Ohanian and Raffo (2012) (OECD OR) dataset. Ohanian and Raffo (2012) present a new, comprehensive dataset on total hours worked at a quarterly frequency, covering 14 OECD countries for the past 50 years. Their dataset is compiled from various international sources and allows us to significantly extend the sample for AU (1970:Q1-2016:Q4) and CA (1961:Q2-2016:Q4). The version of Ohanian and Raffo (2012)'s dataset, which we are working with ends in 2016:Q4. In Appendix 2.B.7 we compare the measure for hours worked between the EO and OR datasets and show that they are very similar for overlapping periods. Additionally, in Appendix 2.B.8, we compare for the US the datasets stemming from the OECD EO dataset, the OECD OR dataset, and the ACD dataset and show that the resulting series from each dataset are also near-identical.

Second, we use annual utilization adjusted TFP series as constructed in Huo et al. (2020)

²See Lequiller and Blades (2014) for details on national accounts.

and Comin, Quintana, et al. (2022) in order to draw conclusions about the disconnect of the MBC shock with TFP.

Third, we use measures of capacity utilization (CU_t) to attribute increases in labor productivity to a pro-cyclical rise in utilization rather than to a technological improvement.

We perform all estimations for the OECD EO and OECD OR datasets with the longest available sample period per country and identical starting dates (1981:Q1; balanced panel or BP henceforth). This allows the use of all the available data and a balanced comparison of results across datasets and countries. Results across datasets and sample lengths are very similar. Unless otherwise specified, the exposition in this paper focuses on the results obtained from the longest available OECD EO dataset per country.

Descriptive Statistics

Below, we provide descriptive statistics to verify that standard unconditional cyclical components of the domestic main macroeconomic variables are similar across the advanced economies in our sample.³ While similarity in unconditional moments is a necessary condition, it is not sufficient to claim that they are also driven by the same type of shock - the central issue investigated in the course of this paper to assert that business-cycles are all alike. To isolate the cyclical fluctuations between 6 and 32 quarters, we use Fitzgerald and Christiano (2003)'s Bandpass-filter. The left panel of Table 2.1 provides a visualization of the comovement over the business-cycle. The right panel quantifies the amplitude of business-cycle fluctuations and the relative volatilities of each variable.

Table 2.1: Correlations and standard deviations (bandpass-filtered, 6-32 quarters)

$\rho(\cdot, Y_t)$	AU	CA	FR	IT	UK	US	σ	AU	CA	FR	IT	UK	US
Y_t	1.00	1.00	1.00	1.00	1.00	1.00	Y_t	1.01	1.30	0.84	1.34	1.06	1.36
I_t	0.89	0.85	0.89	0.91	0.85	0.94	I_t	5.46	4.04	3.87	3.69	4.93	3.80
C_t	0.20	0.84	0.56	0.73	0.78	0.89	C_t	0.87	0.63	0.80	0.77	0.97	0.82
h_t	0.63	0.84	0.27	0.53	0.54	0.87	h_t	1.19	0.78	0.91	0.51	0.98	0.87
u_t	-0.71	-0.88	-0.61	-0.47	-0.63	-0.88	u_t	0.60	0.46	0.37	0.26	0.40	0.49

Contemporaneous correlation with output.

In absolute terms for output and relative to output for remaining variables.

Overall, the unconditional variance-covariance matrices of domestic variables are quite similar across countries. Investment, consumption, hours worked, and employment exhibit relatively strong comovement with output. The only exceptions are consumption in Australia and hours

³Descriptive statistics for all variables are provided in Table 2.16 in Appendix 2.B.5.

worked in France. Regarding the overall business-cycle volatility, we see that the US is on the more volatile end while France has the least volatile business-cycle. In terms of relative volatilities, the countries are again quite similar. Investment is more volatile than output, and consumption is less volatile than output. Hours worked are almost as volatile as output. The only exception is Italy, where hours worked are considerably less volatile than output. For Canada, we find that hours worked are significantly more volatile in the longer OECD OR dataset with a relative standard deviation of 0.95 instead of 0.78 as in the OECD EO dataset. In terms of labor market variables, the moments for hours and unemployment divide the considered countries into two groups: the Anglo-Saxon countries (AU, CA, UK, US) and the Latin countries (FR, IT). In line with Ohanian and Raffo (2012), who establish that the relative volatility and the procyclicality of hours in Euro countries are lower than in Canada, the UK, and the US, we observe that France has the lowest comovement between hours and output, and Italy has the lowest relative volatility of hours worked. These differences are consistent with the view that cross-country variations in employment protection legislation affect how firms adjust labor input. According to the OECD's employment protection ranking, Italy and France have some of the highest levels of protection. In contrast, the UK, Australia, Canada, and the US have some of the lowest levels.⁴ Despite the idiosyncratic differences in labor market variables, the overall conformity in unconditional moment supports the notion that business-cycles are all alike. The remainder of this paper thus seeks to provide evidence that macroeconomic fluctuations in different countries are also triggered by the same type of shock.

2.3.2 Estimation

We estimate VAR and VECM models. Estimations are performed with Bayesian methods, using a Minnesota prior. The posterior distributions are obtained with Gibbs sampling using 50,000 draws. The reported highest posterior density intervals (HPDI) are obtained by the approach described in Koop (2003). The VAR takes the following form:

$$X_t = A_0 + A_1 X_{t-1} + A_2 X_{t-2} + \dots + A_p X_{t-p} + \nu_t, \quad (2.3.1)$$

where $X_t = [X_{1,t}, \dots, X_{N,t}]'$ is the $N \times 1$ vector containing the macroeconomic series observed at time t , p denotes the number of lags⁵, and ν_t is the vector of residuals with covariance matrix

⁴Available on the OECD's [webpage](#). Using a panel of 26 countries with complete data from 1990 to 2019, we find that Italy and France rank 4th and 8th, respectively, while the UK, Australia, Canada, and the US rank 22nd, 23rd, 25th, and 26th, respectively.

⁵We use $p=2$, which is the same as in ACD and confirmed by standard information criteria (Schwartz and Hannan-Quinn).

$E[\nu_t \nu_t'] = \Sigma$. The VECM representation is specified as follows:

$$\Delta X_t = \Psi_0 \Lambda X_{t-1} + \sum_{i=1}^p \Psi_i \Delta X_{t-i} + \nu_t, \quad (2.3.2)$$

where Λ is the matrix of cointegration coefficients, Ψ_0 is the matrix of loadings of these cointegration relationships, Ψ_i are the dynamic coefficients, and the rest of the notation is the same as in the VAR representation in Equation (2.3.1). Following ACD, we estimate two VECMs that differ with respect to the number of unit roots and cointegration relations. In the first case (VECM1), and for the baseline estimation - in which only domestic variables are considered - we assume that the real quantities (Y_t, I_t, C_t) and LP_t share a single stochastic trend. In contrast, the remaining variables $y \in \{h, u, R, \pi\}$ are assumed to be stationary. The cointegrating relationship is of the type $x_t = \alpha_x + \beta_x LP_t$ for $x \in \{Y, I, C\}$. In the second case (VECM2), we again assume that the real quantities (Y_t, I_t, C_t) and LP_t share a single stochastic trend and additionally consider that the nominal variables (π_t, R_t) share another stochastic trend, while the remaining variables (u_t, h_t) are assumed to be stationary. The resulting cointegrating relationships are of the type $x_t = \alpha_x + \beta_x LP_t$ for $x \in \{Y, I, C\}$ and $R_t = \delta + \gamma \pi_t$.⁶ As in ACD, we assume that the long-run growth of real quantities is driven by technology. In the VECM2, nominal variables are assumed to be driven by nominal growth, which is disconnected from growth in real variables.

2.3.3 Identification

In what follows, we outline ACD's identification in more detail making the present paper self-contained. First, for each country separately, we estimate a VAR or VECM and obtain the reduced form residuals ν_t . Let us map these residuals ν_t into the underlying unit-variance shocks ϵ_t as follows: $\nu_t = S\epsilon_t \Leftrightarrow \epsilon_t = S^{-1}\nu_t$. The structural impact matrix S can be further decomposed into $S = \tilde{S}Q$, where \tilde{S} is the Cholesky decomposition of the covariance matrix Σ of the residuals ν_t . Note that Q is an orthonormal matrix such that $Q^{-1} = Q'$ and satisfies $QQ' = I$ by construction. S is an invertible $N \times N$ matrix and satisfies $SS' = \Sigma$. ϵ_t is i.i.d. over time with $E(\epsilon_t \epsilon_t') = I$. We can thus rewrite the shocks as follows:

$$\nu_t = \tilde{S}Q\epsilon_t \Leftrightarrow \epsilon_t = Q'\tilde{S}^{-1}\nu_t. \quad (2.3.3)$$

Note that a shock $\epsilon_t^{[j]}$ in ϵ_t corresponds to the column $q^{[j]}$ in Q . Second, let the Wold representation of the VAR in (2.3.1) be $X_t = A(L)^{-1}\nu_t = B(L)\nu_t$, and combine it with (2.3.3) to

⁶For the estimation including open economy variables, TOT are added to y . EXP and IMP are added to the set x .

get⁷

$$X_t = A(L)^{-1}\nu_t = B(L)\nu_t = B(L)\tilde{S}Q\epsilon_t = C(L)Q\epsilon_t = \Gamma(L)\epsilon_t,$$

where $C_\tau \equiv B_\tau\tilde{S}$ and $\Gamma_\tau = C_\tau Q$. The IRFs of the variables in X_t to a shock in ϵ_t are encompassed in $\{\Gamma\}_{\tau=0}^\infty$, which we obtain from the sequence $\{C\}_{\tau=0}^\infty$ that encapsulates the Cholesky transformation of the VAR residuals. Third, consider three distinct objects: k , the k th variable in X_t ; j , the j th shock in ϵ_t ; and q , the j th column in Q . The effect of a shock j on the variable k at horizon τ is then given by the (k, j) element of $\Gamma_\tau = C_\tau Q$. Since we perform identification in the frequency rather than the time domain, we denote the contribution of shock j to the *spectral density* of variable k over the frequency band $[\underline{\omega}, \bar{\omega}]$ by:

$$\begin{aligned} \Upsilon(q; k, \underline{\omega}, \bar{\omega}) &\equiv \int_{\omega \in [\underline{\omega}, \bar{\omega}]} \left(\overline{C^{[k]}(e^{-i\omega})} q C^{[k]}(e^{-i\omega}) q \right) d\omega \\ &= q' \left(\int_{\omega \in [\underline{\omega}, \bar{\omega}]} \overline{C^{[k]}(e^{-i\omega})} C^{[k]}(e^{-i\omega}) d\omega \right) q \\ &= q' \Theta(k, \underline{\omega}, \bar{\omega}) q, \end{aligned} \tag{2.3.4}$$

where \bar{u} denotes the complex conjugate transpose of any vector u and $C_\tau^{[k]}$ denotes the k th row of matrix C_τ . The last expression in (2.3.4) determines the contribution of any structural shock to the matrix $\Theta(k, \underline{\omega}, \bar{\omega})$, which is computed from the data and represents the total variability of variable k over the frequency band $[\underline{\omega}, \bar{\omega}]$ in terms of the contributions of all the Cholesky-transformed residuals. Fourth, to identify the shock as the shock that maximizes the contribution to the volatility of a particular variable over a specific frequency band, it suffices to choose q to maximize (2.3.4). It follows that q is the eigenvector corresponding to the largest eigenvalue of the matrix $\Theta(k, \underline{\omega}, \bar{\omega})$. By varying the targeted variable k in X_t and the frequency band $[\underline{\omega}, \bar{\omega}]$, we assemble a collection of shocks. Notably, the identified shocks in this collection are not necessarily orthogonal or restricted in any way to be orthogonal to one another. Each shock is identified such that it explains most of the volatility of a particular variable over a specific frequency band. By targeting different variables, the same shock may be identified for multiple variables. Regarding the frequency bands, we use two frequency bands; one typically associated with the business-cycle (6-32 quarters), where $[\underline{\omega}, \bar{\omega}] = [2\pi/32, 2\pi/6]$, and the other one associated with the long-run (80- ∞ quarters), where $[\underline{\omega}, \bar{\omega}] = [0, 2\pi/80]$. The first core subset of our anatomy consists of the shocks that target the key macroeconomic variables (unemployment, output, hours worked, consumption and investment) over the business-cycle frequencies. The second core subset comprises the shocks that target output, consumption,

⁷We use the fact that $A(L) = \sum_{\tau=0}^p A_\tau L^\tau$, $B(L) = \sum_{\tau=0}^p B_\tau L^\tau$, $C(L) = \sum_{\tau=0}^p C_\tau L^\tau$, $\Gamma(L) = \sum_{\tau=0}^p \Gamma_\tau L^\tau$, with L denoting the backshift operator and $A(0) = A_0 = I$.

investment, and labor productivity over the long-run frequencies. Finally, we consider the properties of the shocks that target the open economy variables (exports, imports, and terms of trade). Together, these subsets build up the *cross-country* anatomy, whose properties are discussed in detail in the following sections. In what follows, we mainly discuss the results obtained from the VECM1 specification, which provide substantial evidence towards a MBC shock in each of the studied countries. When estimating the VAR, the evidence towards a MBC shock is weaker but aligns with the results obtained in Miyamoto et al. (2022), who establish the existence of a MBC shock in each of the G7 countries using a VAR. We thus delegate the results of the VAR specification to the Appendix, Section 2.C.2 for the closed economy, and Section 2.C.4 for the open economy version. Finally, because the results for the VECM2 are in line with those obtained with the VECM1, they are omitted.

2.4 Cross-Country Analysis

In this section, we report the results of the baseline estimation, which includes eight domestic variables: unemployment, output, hours worked, investment, consumption, labor productivity, inflation, and the nominal interest rate. As a first step, we focus on the first core subset of shocks, which are the shocks targeting the key macroeconomic variables at business-cycle frequencies. We characterize the statistical properties of these shocks and how they relate to the nominal side, technology, and the long-run. In the second step, we turn to the shocks that target output, consumption, investment, and labor productivity at long-run frequencies. We report how they relate to the long-run and quantify their contributions at business-cycle frequencies. Finally, we investigate how strongly linked the identified main business-cycle drivers are internationally. To foreshadow, the evidence presented below asserts that business-cycles in advanced economies are driven by the same type of non-inflationary aggregate demand force that is independent of technology and the long-run and plays a significant role in international business-cycle comovement.

2.4.1 Existence of a Main Business-Cycle Shock

To establish the existence of an MBC shock, two key pieces of evidence are required. The first is that the IRFs of the individual shocks share the same dynamics across variables, as shocks that are observationally equivalent in terms of their IRFs are essentially the same shock. The second piece of evidence is that the shocks account not only for the bulk of business-cycle volatility in the variable they target but also for a large part of the volatility in the other key macroeconomic variables.

Let us consider the first requirement. Figure 2.1 reports for every country the IRF of each

variable in response to the shocks targeting unemployment, output, investment, hours, and consumption at business-cycle frequencies.⁸ The overall evidence derived from the IRFs highlights that in all countries, the effects of the shocks that target *any* of the key real macroeconomic quantities are highly interchangeable. The only exceptions are the consumption shock in Australia, which shows a lower connection, and the hours shock in France, which is less tightly connected to the other facets but nevertheless exhibits very similar dynamics. The overall picture in Figure 2.1 drives home the point that within each economy, any of the identified shocks is essentially a facet of the same shock, the MBC shock. What are the dynamic cross-variable properties of this shock? We observe that the shock triggers transitory comovement in the key macroeconomic variables and *peaks* after about one year. Specifically, we observe that the shock *decreases* unemployment, and *increases* output, hours worked, investment, and consumption. The shock moreover triggers an initially small positive effect in labor productivity, which either becomes statistically insignificant or negative quickly. Regarding the nominal side, the shock causes inflation to increase only slightly and triggers a pronounced transitory *increase* in the nominal interest rate. The IRFs underpin the notion of a single business-cycle force within each of the advanced economies considered. Reassuringly, the observed dynamics are also in line with the results found in ACD for the US and in Miyamoto et al. (2022), who use a VAR instead of a VECM and variables in relative terms to the rest of the world.

With regard to the second necessary requirement to make the case for an MBC shock, let us consider Table 2.2. This table quantifies the amount of business-cycle volatility each shock explains for each variable.⁹ These metrics allow us to quantitatively interpret the footprint of each shock on the core macroeconomic variables in each country. We can distinguish two subtypes of metrics that are of special interest. The first sub-metric includes the five diagonal elements in the left part of each sub-table. These numbers tell us how much of the volatility any given shock explains of the variable that it targeted. We observe that these metrics are uniformly large, varying between 60% and 80%.

The second crucial metric encompasses the off-diagonal elements. These elements tell us how much a shock explains of the business-cycle volatility of the variables *other* than its target. To be considered as the main business-cycle shock, these metrics ought to be close to or above 50%. Again, the results paint a consistent pattern across countries, with the majority of the off-diagonal volatilities being close to or above the threshold. Some country-specific differences, as already reported in the descriptive statistics and IRFs, stand out.

⁸For IRFs, we display the median of the HPDI. The shaded area denotes the 68% HPDI for the unemployment shock.

⁹For reasons of space, we provide only the median of the posterior distribution for variance contributions. Tables including the 68% HPDI are available upon request.

Figure 2.1: The various facets of the MBC shock, IRFs (VECM1)




Table 2.2: The various facets of the MBC shock, variance contributions (VECM 1)

	u	Y	h	I	C	LP	π	R
AU								
Unemployment	83.6	60.9	68.6	67.0	23.3	28.4	21.5	46.3
Output	60.2	81.4	47.6	62.2	35.6	39.2	8.9	29.7
Hours Worked	66.9	46.4	83.9	51.3	14.7	21.7	15.9	34.4
Investment	64.6	60.9	51.6	85.8	25.0	29.0	11.7	41.0
Consumption	15.9	20.6	12.2	12.5	62.8	12.1	5.5	17.9
CA								
Unemployment	87.5	72.0	81.1	60.3	66.6	28.8	16.8	47.9
Output	74.4	82.9	74.4	64.1	64.8	41.5	20.8	48.7
Hours Worked	81.6	72.2	87.6	58.3	65.2	27.9	13.6	45.1
Investment	63.1	63.8	60.2	83.2	57.9	29.9	24.2	39.1
Consumption	68.0	64.5	66.6	57.5	85.5	26.3	12.0	38.3
FR								
Unemployment	68.6	42.8	22.4	45.3	26.7	29.3	13.6	15.7
Output	46.2	67.0	13.3	52.7	53.9	63.3	15.5	15.8
Hours Worked	22.1	11.8	75.4	15.0	11.5	7.1	9.5	10.8
Investment	51.2	54.9	16.9	65.4	38.4	45.2	15.3	15.2
Consumption	24.2	45.8	11.9	31.7	77.8	45.7	7.4	8.3
IT								
Unemployment	80.4	34.0	43.0	41.8	31.4	15.9	10.2	11.7
Output	42.3	67.0	44.1	59.2	42.1	49.0	9.5	23.6
Hours Worked	46.0	37.0	81.5	36.9	32.6	14.0	10.6	11.6
Investment	44.8	54.1	39.2	74.9	34.7	40.1	5.3	21.9
Consumption	37.0	41.9	37.4	36.9	68.5	24.3	6.2	15.4
UK								
Unemployment	80.8	68.9	70.6	70.5	63.3	33.0	18.5	48.6
Output	63.3	86.1	53.8	64.8	74.0	57.8	18.5	33.8
Hours Worked	71.7	59.8	80.1	69.9	55.3	25.9	17.4	40.4
Investment	68.4	68.1	67.2	81.2	61.7	36.5	18.8	43.3
Consumption	57.2	73.8	47.9	57.8	86.0	46.3	12.6	26.8
US								
Unemployment	82.1	65.3	78.0	65.3	53.2	27.7	13.2	31.8
Output	69.0	78.0	69.3	68.7	65.5	36.5	7.8	25.3
Hours Worked	75.6	63.7	84.5	60.2	55.6	26.9	11.1	28.4
Investment	68.9	68.7	65.1	77.3	52.7	32.2	9.8	31.0
Consumption	56.5	66.6	61.1	54.6	77.8	29.2	7.1	16.2

The rows correspond to different targets in the construction of the shock. The columns give the contributions of the constructed shock to the business-cycle volatility of the variables.

Consider first the US, Canada, and the UK. There, the main shocks explain close to or at least half of the business-cycle volatility in their respective off-diagonal elements in unemployment, output, hours, investment, and consumption. For Australia, the same pattern emerges with the exception of the off-diagonal elements of consumption. For France, the shock that targets hours worked explains little of any variable and the off-diagonal elements with respect to hours worked are also very low. The other off-diagonal elements are, however, in the ballpark of the other countries. For Italy, the key shocks explain slightly less than half of the volatilities of the variables other than their respective target. In line with the observations about the differences in labor market variables in Section 2.3.1, for France and Italy we thus observe similar dynamics but a weaker connection between the main business-cycle shocks on the good and the labor market when compared to the Anglo-Saxon countries.

Finally, we note that none of the shocks in any of the countries accounts for a significant fraction of the business-cycle volatility of inflation.

Interchangeability

To pin down the interchangeability of the key facets of the MBC shock in succinct quantitative terms, we consider conditional time series. Conditional time series are constructed by predicting the value of a variable as produced by a particular shock alone. Relative to pure innovations, conditional time series include information on how the innovations propagate over time. Table 2.3 depicts for each variable and each country, the average of the bandpass-filtered correlations between the conditional time series as produced by the output shock and that generated by the investment, consumption, hours worked and unemployment shock. The uniformly strong correlations underscore the strong interchangeability among key facets of the MBC shock as observed in the IRFs.¹⁰

Table 2.3: Average correlation of conditional time series (VECM1)

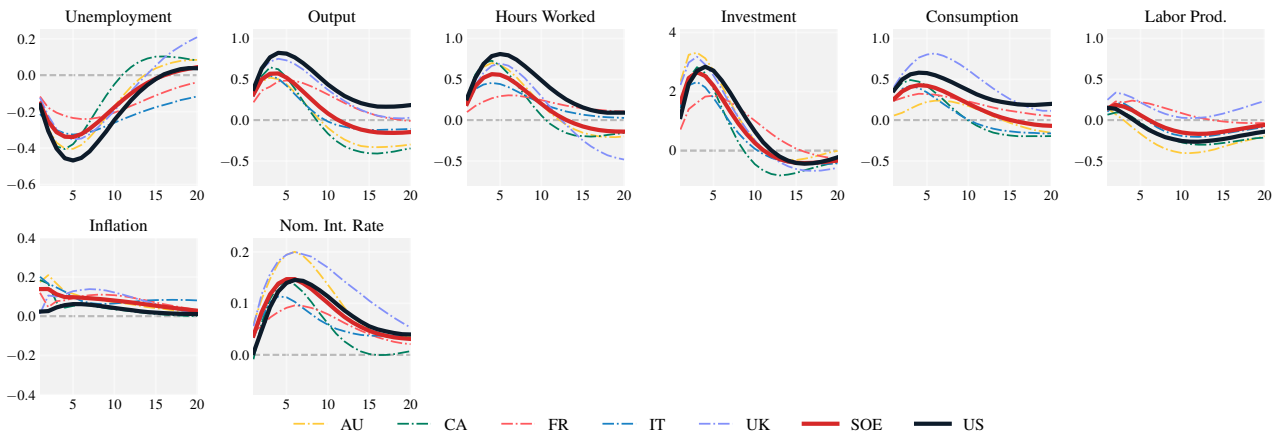
	AU	CA	FR	IT	UK	US
Unemployment	0.76	0.95	0.78	0.73	0.92	0.92
Output	0.94	0.98	0.76	0.90	0.98	0.97
Investment	0.85	0.96	0.77	0.86	0.92	0.96
Consumption	0.67	0.93	0.87	0.81	0.96	0.95
Hours Worked	0.73	0.96	0.91	0.83	0.88	0.92

¹⁰The correlations between each series and output are provided in Table 2.20 in Appendix 2.C.1.

MBC shock around the World

The above evidence highlights the interchangeability of the shocks that target any of the main macroeconomic quantities over business-cycle frequencies. These shocks produce nearly indistinguishable IRFs, are highly correlated, and each shock accounts not only for the bulk of the volatility of its target but also for most of the volatility of the other main variables. As in ACD, we conclude that these shocks are essentially all facets of a single common force (the MBC shock) that explains the majority of business-cycle fluctuations in each country. Let us now shed some light on the similarity of the identified MBC shocks across countries. Figure 2.2 compares the shocks that target unemployment. Because the shocks that constitute the MBC shock are interchangeable, the specific choice of the target does not alter the picture. We plot a summary IRF calculated from averaging the IRF of AU, CA, FR, IT, and the UK, denoted SOE (solid red), and compare it to the US (solid black) to provide a comprehensive overview.

Figure 2.2: IRFs to the MBC shocks around the world (VECM1)



IRFs of all the variables to the MBC shock (identified as the shock that targets unemployment). Horizontal axis: Time horizon in quarters. Vertical axis: Percent.

Consistent with the previous results, we find that the dynamics of the MBC shocks in SOE countries and the US are highly alike and characterized by a strong, positive, and transitory comovement in unemployment, output, hours, investment, and consumption over the business-cycle. Moreover, the MBC shock causes inflation to only slightly increase while it triggers a pronounced transitory increase in the nominal interest rate. While our comparison underpins the similarity of the shock across countries, it also highlights some interesting country-specific behavior in magnitude and persistence: First, compared to the other countries, consumption in Australia reacts only very little and peaks last. Second, the reactions of unemployment and hours worked are the weakest in France. Compared to the US, unemployment reacts only about half as much. Third, the reaction of consumption is strongest for the UK, peaking at around 5 quarters. Fourth, unemployment is most persistent in Italy, where it takes about twice as long

to revert.

The properties of the identified MBC shocks are robust across datasets and sample lengths. Table 2.4 provides the variance contributions of the MBC shocks across the OECD EO and OECD OR datasets as well as their respective BP versions, starting in 1981:Q1. The MBC shocks explains the overwhelming majority of the volatility in unemployment and about 50-60% of the volatility of output, hours, and investment at business-cycle frequencies. In line with previous observations, the explained volatility is somewhat lower for consumption. Finally, the footprint on inflation is uniformly low.

Table 2.4: Variance contribution of the MBC shocks around the world (VECM1)

	u	Y	h	I	C	LP	π	R
EO								
SOE	80.2	55.7	57.1	57.0	42.2	27.1	16.1	34.0
US	82.1	65.3	78.0	65.3	53.2	27.7	13.2	31.8
OR								
SOE	79.8	46.5	50.9	53.7	33.9	19.3	18.3	33.2
US	76.3	53.7	67.3	60.0	40.3	20.9	16.1	30.9
EO BP								
SOE	81.1	57.8	59.6	60.4	47.4	26.4	13.2	38.1
US	85.0	72.9	81.5	73.5	69.7	37.2	14.7	44.0
OR BP								
SOE	81.0	54.6	60.2	63.3	47.2	25.0	16.2	39.2
US	85.5	72.2	84.3	71.6	68.4	41.8	17.7	46.6

Columns give the contributions of the MBC shock (identified as the shock that targets unemployment) to the business-cycle volatility of the variables. EO refers to the dataset using data uniquely from the OECD EO database. OR refers to the dataset that uses data from Ohanian and Raffo (2012) for hours, employment and population. BP refers to the balanced panel version of each dataset, with identical starting dates (1981:Q1).

Having established the existence of a main business-cycle shock in each of the studied countries, we now examine the characteristics of the identified MBC shocks in more detail. We will analyze how the MBC shocks are related to the nominal side, technology, and the long-run. Examining these properties is vital for determining the *type* of shock the MBC shock represents.

2.4.2 Nominal side

Table 2.5 describes the relationship between real economic activity and inflation and complements the evidence seen in Table 2.2. In all countries, the identified MBC shocks account only for very little of the business-cycle volatility in inflation. Conversely, the shock that targets inflation explains the vast majority of the business-cycle volatility in inflation but explains only

very little of either unemployment or output. These results replicate the weak link between the force that accounts for the bulk of real economic activity at business-cycle frequencies and the main driver of inflation as documented in ACD for the US as well as in Miyamoto et al. (2022) for the G7.

Table 2.5: Inflation and the business-cycle (VECM1)

	u	Y	π	R		u	Y	π	R
<hr/>					<hr/>				
AU					IT				
Unemployment	83.6	60.9	21.5	46.3	Unemployment	80.4	34.0	10.2	11.7
Output	60.2	81.4	8.9	29.7	Output	42.3	67.0	9.5	23.6
Inflation	12.6	6.0	74.9	15.2	Inflation	9.3	7.0	87.4	32.6
Nom. Int. Rate	54.3	34.8	24.6	65.2	Nom. Int. Rate	15.0	29.1	33.0	89.3
<hr/>					<hr/>				
CA					UK				
Unemployment	87.5	72.0	16.8	47.9	Unemployment	80.8	68.9	18.5	48.6
Output	74.4	82.9	20.8	48.7	Output	63.3	86.1	18.5	33.8
Inflation	9.9	15.5	83.1	7.6	Inflation	5.7	6.8	68.1	7.0
Nom. Int. Rate	49.2	46.5	13.7	79.7	Nom. Int. Rate	47.3	32.1	24.6	81.8
<hr/>					<hr/>				
FR					US				
Unemployment	68.6	42.8	13.6	15.7	Unemployment	82.1	65.3	13.2	31.8
Output	46.2	67.0	15.5	15.8	Output	69.0	78.0	7.8	25.3
Inflation	8.3	10.0	74.8	22.6	Inflation	11.4	10.9	87.4	22.7
Nom. Int. Rate	18.0	21.3	24.4	78.8	Nom. Int. Rate	26.4	23.5	21.7	68.2
<hr/>					<hr/>				

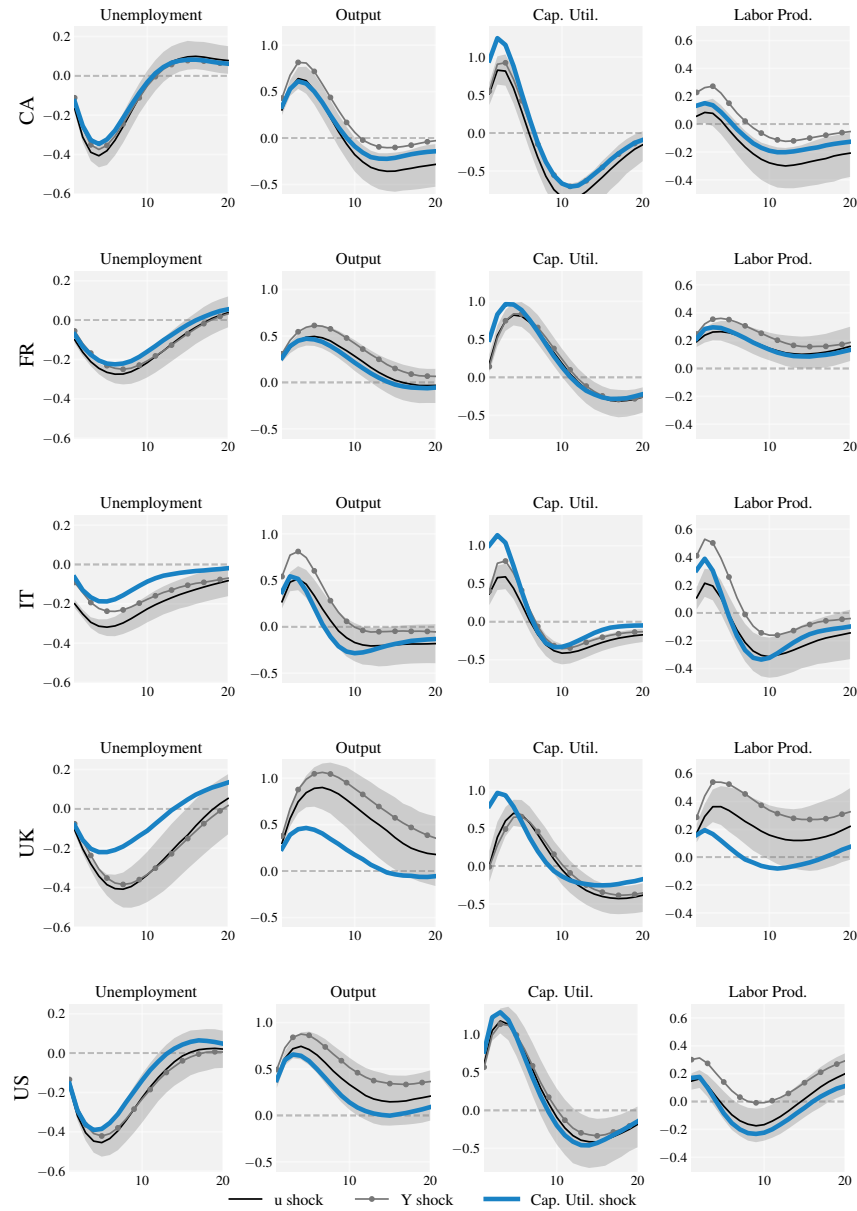
The rows correspond to different targets in the construction of the shock. The columns give the contributions of the constructed shock to the business-cycle volatility of the variables.

2.4.3 Technology

Consider the connection between the MBC shock and technology. As quarterly, long-horizon utilization-adjusted TFP series are only available for the US, the relationship between the MBC shock and TFP cannot be directly observed for the SOE. We resort to two pieces of indirect evidence. First, we calculate the correlation between the annualized MBC shocks and *yearly* utilization-adjusted TFP series as constructed in Huo et al. (2020) and Comin, Quintana, et al. (2022). The two approaches differ in how they adjust TFP from unobserved variation in capacity utilization rates over the business-cycle. Huo et al. (2020) use hours per worker as a proxy for capacity utilization, and Comin, Quintana, et al. (2022) rely on capacity utilization surveys. For the SOE, we find only very low average correlations in the ballpark of 0.2 for the MBC shocks as constructed by targeting unemployment. This is in line with ACD's observation that

the US MBC shock is disconnected from technology.

Figure 2.3: The MBC shock and capacity utilization, IRFs (VECM1)



Horizontal axis: Time horizon in quarters. Vertical axis: Percent.

Second, we amend the baseline dataset with a measure for capacity utilization for Canada, France, Italy, the UK, and the US. For Australia, the ACCI-Westpac Survey of Industrial Trends reports capacity utilization in manufacturing, but we could not obtain the series from the provider. Including a measure for capacity utilization allows us to compare the facets of the MBC shock with the shock that targets capacity utilization at business-cycle frequencies. This in turn helps us attribute the increase in labor productivity to an increase in capacity utilization rather than an increase in technology. Figure 2.3 provides the corresponding IRF

and Table 2.6 the corresponding variance contributions.

Table 2.6: The MBC shock and capacity utilization, variance contributions (VECM 1)

	u	Y	CU	LP
CA				
Unemployment	85.0	68.0	61.5	27.1
Output	70.5	80.1	57.8	40.1
Cap. Util.	62.3	57.5	80.6	26.1
FR				
Unemployment	77.1	60.7	59.0	42.0
Output	62.3	76.3	59.7	66.8
Cap. Util.	58.5	57.9	73.4	51.6
IT				
Unemployment	74.6	32.8	33.6	20.5
Output	40.6	62.2	47.1	47.8
Cap. Util.	31.5	39.7	76.5	37.7
UK				
Unemployment	78.9	70.7	48.2	38.5
Output	66.1	85.2	41.7	61.7
Cap. Util.	43.7	36.4	65.5	24.5
US				
Unemployment	72.2	53.5	55.7	33.8
Output	60.6	64.6	51.1	37.0
Cap. Util.	61.5	50.9	64.6	42.1

The rows correspond to different targets in the construction of the shock. The columns give the contributions of the constructed shock to the business-cycle volatility of the variables. CU stands for capacity utilization.

The properties of the MBC shock do not change when capacity utilization is included.¹¹ The MBC shock and the shock that target capacity utilization are closely connected and nearly identical in terms of magnitude and dynamics. Both shocks increase capacity utilization as well as labor productivity. Only in the UK is the magnitude of the capacity utilization shock slightly lower than that of the MBC shock. Altogether, the above evidence attributes the observed transitory increases in labor productivity to the pro-cyclical increase in utilization rather than to a true technological improvement. To support this conclusion, we expand ACD's dataset for the US to include capacity utilization. As displayed in Figure 2.12 in Appendix 2.C.1, the MBC shock and the capacity utilization shock again lead to increases in capacity utilization and labor productivity, whereas TFP does not increase.

¹¹The full set of IRFs and corresponding variance contributions are provided in Appendix 2.C.1, Figure 2.12 and Table 2.21.

2.4.4 Long-Run

Having investigated the characteristics of the MBC shock with regard to inflation and technology at business-cycle frequencies, we now turn to the corresponding footprint of the MBC shock at long-run frequencies. Tables 2.7 and 2.8 characterize the variance decompositions of the business-cycle shocks in the long-run.¹² As in ACD, we also discover a pronounced disconnect between the short- and long-run contributions. The footprint of the short-run shocks are negligible in the long-run. This echoes ACD's conclusion that the main driver of the business-cycle is distinct from the factors influencing long-run productivity and output.

Table 2.7: MBC shock, variance contributions at long-run frequencies ($80-\infty$ Q) (VECM 1)

	u	Y	h	I	C	LP	π	R
AU								
Unemployment	14.5	5.7	14.3	5.7	5.7	5.7	24.5	24.1
Output	10.1	10.4	8.9	10.4	10.4	10.4	12.0	10.8
Hours Worked	8.2	3.4	13.3	3.4	3.4	3.4	8.6	9.3
Investment	13.2	8.7	12.6	8.7	8.7	8.7	15.4	21.1
Consumption	7.6	4.8	5.6	4.8	4.8	4.8	6.1	6.0
CA								
Unemployment	8.6	4.1	11.1	4.1	4.1	4.1	21.7	11.6
Output	7.4	2.7	11.5	2.7	2.7	2.7	17.9	10.5
Hours Worked	9.0	6.4	14.0	6.4	6.4	6.4	15.1	10.6
Investment	5.3	2.2	6.8	2.2	2.2	2.2	16.9	7.9
Consumption	7.8	2.7	11.2	2.7	2.7	2.7	16.0	9.8
FR								
Unemployment	6.9	2.9	3.8	2.9	2.9	2.9	12.9	10.2
Output	5.2	2.8	3.6	2.8	2.8	2.8	11.9	6.5
Hours Worked	6.1	4.3	8.1	4.3	4.3	4.3	16.7	24.5
Investment	6.9	4.5	4.5	4.5	4.5	4.5	13.9	6.9
Consumption	5.0	3.2	4.7	3.2	3.2	3.2	7.0	5.4

The rows correspond to different targets in the construction of the shock. The columns give the contributions of the constructed shock to the long-run volatility of the variables.

¹²For the US, we use the OECD OR dataset. Relative to the OECD EO dataset, its long-run properties are more in line with ACD.

Table 2.8: MBC shock, variance contributions at long-run frequencies ($80\text{-}\infty$ Q) (VECM 1)

	u	Y	h	I	C	LP	π	R
IT								
Unemployment	23.4	8.0	8.4	8.0	8.0	8.0	20.5	10.9
Output	13.5	8.0	8.6	8.0	8.0	8.0	13.5	11.1
Hours Worked	27.8	15.7	26.5	15.7	15.7	15.7	27.7	24.7
Investment	7.8	2.2	3.2	2.2	2.2	2.2	4.9	4.0
Consumption	10.3	7.7	5.4	7.7	7.7	7.7	8.6	7.1
UK								
Unemployment	13.8	23.7	15.2	23.7	23.7	23.7	16.8	22.3
Output	14.1	33.5	13.0	33.5	33.5	33.5	15.2	19.0
Hours Worked	14.2	20.4	15.8	20.4	20.4	20.4	20.0	24.3
Investment	14.0	25.6	13.8	25.6	25.6	25.6	16.1	23.8
Consumption	9.1	14.7	7.5	14.7	14.7	14.7	6.6	9.1
US (OR)								
Unemployment	24.6	5.6	24.4	5.6	5.6	5.6	16.7	19.1
Output	22.6	3.6	20.6	3.6	3.6	3.6	10.7	11.9
Hours Worked	19.2	8.0	24.9	8.0	8.0	8.0	14.7	15.0
Investment	20.5	3.6	17.8	3.6	3.6	3.6	10.5	14.6
Consumption	29.0	6.9	18.2	6.9	6.9	6.9	9.2	8.8

The rows correspond to different targets in the construction of the shock. The columns give the contributions of the constructed shock to the long-run volatility of the variables.

Summarizing all of the above, we find that the identified MBC shocks in Australia, Canada, France, Italy, the UK, and the US are highly similar and can be characterized as non-inflationary, demand-side type of shocks, which are disconnected from the forces that drive the economy in the long-run.

2.4.5 Existence of a Main Long-Run Shock

After focusing on the shocks that target the main macroeconomic variables at business-cycle frequencies, we consider the shocks that target output, investment, consumption, and labor productivity at long-run frequencies ($80\text{-}\infty$ quarters).

The IRFs to the long-run shocks are displayed in Figure 2.4.¹³ For all countries, the long-run shocks are nearly indistinguishable from one another and they overwhelmingly exhibit a permanent effect on output, investment, consumption, and labor productivity. The results for Canada, France, Italy, and the US are especially strong. For Australia, the consumption shock is slightly less tightly connected to the other shocks. Finally, while the IRFs for the UK also

¹³We display the median of the HPDI. The shaded area denotes the 68% HPDI for the output shock.

exhibit strong interchangeability, the IRFs are not fully consistent with a permanent effect. The long-run results for the UK are hence the least clear-cut in the VECM1 estimation but improve with the VAR, where they are fully consistent with the pattern of a permanent shock.¹⁴

Table 2.9: Long-run shocks, contributions at long-run frequencies ($80\text{-}\infty$ Q) (VECM 1)

	AU	CA	FR	IT	UK	US (OR)
Output	96.5	98.2	99.8	99.6	99.7	99.0
Investment	88.5	96.7	98.9	98.3	99.2	98.7
Consumption	94.6	99.3	99.8	99.6	99.7	99.1
Labor Productivity	96.8	98.9	99.8	99.6	99.8	99.1

The rows correspond to different targets in the construction of the shock. The columns give the contributions of the constructed shock to the long-run volatility of the variables.

The variance contributions displayed in Table 2.9 complement the interchangeability result from the long-run IRFs. The long-run shocks explain nearly all of the long-run volatility in output, investment, consumption and labor productivity, which lends additional support to the existence of a ‘main long-run’ (MLR) shock. To further characterize the MLR shock, consider Table 2.10. This table shows for each variable the contribution at business-cycle frequencies of the aforementioned long-run shocks. The table quantifies the size of the footprint of the long-run shocks in the short-run and is hence inversely related to Tables 2.7 and 2.8, where the long-run contributions of the short-run shocks were investigated. Consistent with the disconnect of the short-run shocks in the long-run, Table 2.10 shows that the long-run shocks also only have a very small impact on the business-cycle volatilities of any variable. As before, the results for the UK are again more pronounced for the VAR estimation. Consistent with the results in ACD for the US we find that a MLR exists that is disconnected from what drives the business-cycle.

¹⁴See Figure 2.15 in Appendix 2.C.2.

Figure 2.4: The various facets of the MBR shock, IRFs (VECM1)




Table 2.10: Long-run shocks, contributions at business-cycle frequencies (VECM 1)

	u	Y	h	I	C	LP	π	R
AU								
Output	9.2	9.6	7.2	9.5	8.9	13.2	9.2	17.5
Investment	11.3	9.0	8.4	10.1	8.9	12.1	10.1	19.4
Consumption	9.5	9.3	7.0	9.5	9.4	13.6	9.8	18.3
Labor Productivity	10.1	10.2	7.2	9.8	9.1	13.7	9.5	18.0
CA								
Output	6.9	6.8	9.7	3.6	5.6	27.9	5.6	7.3
Investment	7.2	6.5	9.8	4.1	6.0	23.9	5.4	8.4
Consumption	6.7	6.4	9.2	3.6	5.5	25.5	5.2	7.5
Labor Productivity	6.9	6.7	9.8	3.6	5.6	27.1	5.5	7.3
FR								
Output	6.3	7.7	9.7	9.3	6.8	6.5	21.6	25.6
Investment	7.0	8.7	9.2	10.7	7.7	7.1	22.6	27.6
Consumption	6.4	7.6	9.6	9.1	7.0	6.4	21.6	25.3
Labor Productivity	6.4	7.8	9.8	9.3	6.9	6.5	21.4	26.1
IT								
Output	7.8	11.1	14.3	5.1	10.5	10.4	16.6	11.2
Investment	7.3	9.5	13.7	5.4	8.6	8.7	19.5	12.9
Consumption	7.9	11.2	14.3	5.1	10.6	10.5	16.5	11.1
Labor Productivity	7.9	11.6	15.3	5.2	10.8	10.2	16.5	10.8
UK								
Output	21.4	26.9	18.5	21.0	12.0	15.8	21.6	35.2
Investment	20.2	24.8	17.4	18.9	10.6	14.5	21.7	34.8
Consumption	21.1	26.6	18.2	20.8	11.6	15.7	21.6	35.1
Labor Productivity	22.0	27.5	19.1	21.8	12.9	16.4	21.3	35.2
US (OR)								
Output	12.5	14.5	13.3	12.5	17.8	18.4	17.5	9.0
Investment	12.8	15.0	13.2	12.9	18.6	19.2	18.3	9.2
Consumption	12.3	14.5	13.1	12.4	17.7	18.7	17.1	8.7
Labor Productivity	12.1	13.9	13.1	11.8	17.2	20.7	18.7	8.8

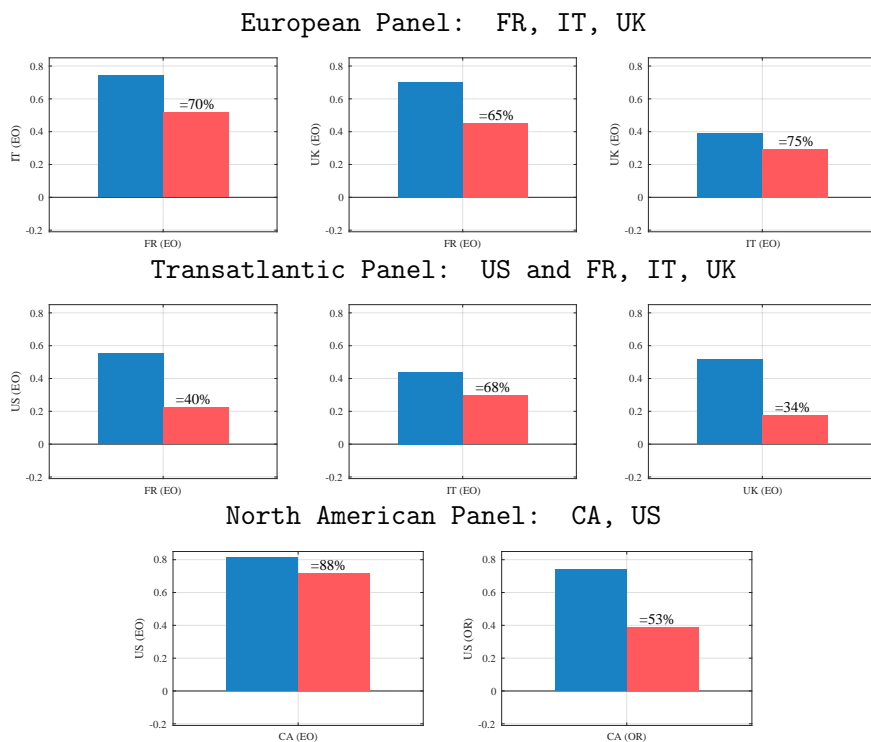
The rows correspond to different targets in the construction of the shock. The columns give the contributions of the constructed shock to the business-cycle volatility of the variables.

2.4.6 International Comovement

The previous sections provided evidence for an MBC shock and a MLR shock in each country considered. We now investigate whether the domestic MBC shocks are linked internationally. We compare the contemporaneous cross-country output correlations at business-cycle frequencies between actual output series (blue bars) and counterfactual output series (red bars). The

latter were generated by isolating the effect of the output shock and excluding all other shocks. In Figure 2.5, we provide the evidence for European (FR, IT, UK), transatlantic (US, FR, IT, UK), and North American (CA, US) country pairs. For the latter, we also provide results for the OECD OR dataset due to the short sample length for Canada in the OECD EO dataset. We do not investigate contemporaneous output comovement with Australia, as its trade is mainly oriented towards the Asia-Pacific region,¹⁵ a region currently not included in our analysis. Percentage numbers on top of the red bars indicate the fraction of counterfactual output correlation to actual output correlation.

Figure 2.5: Cross-country correlation of output: unconditional and conditional (VECM1)



Blue bars indicate the contemporaneous correlation of the bandpass-filtered series for output. Red bars indicate the contemporaneous correlation of bandpass-filtered series for counterfactual output (output as generated by the output shock).

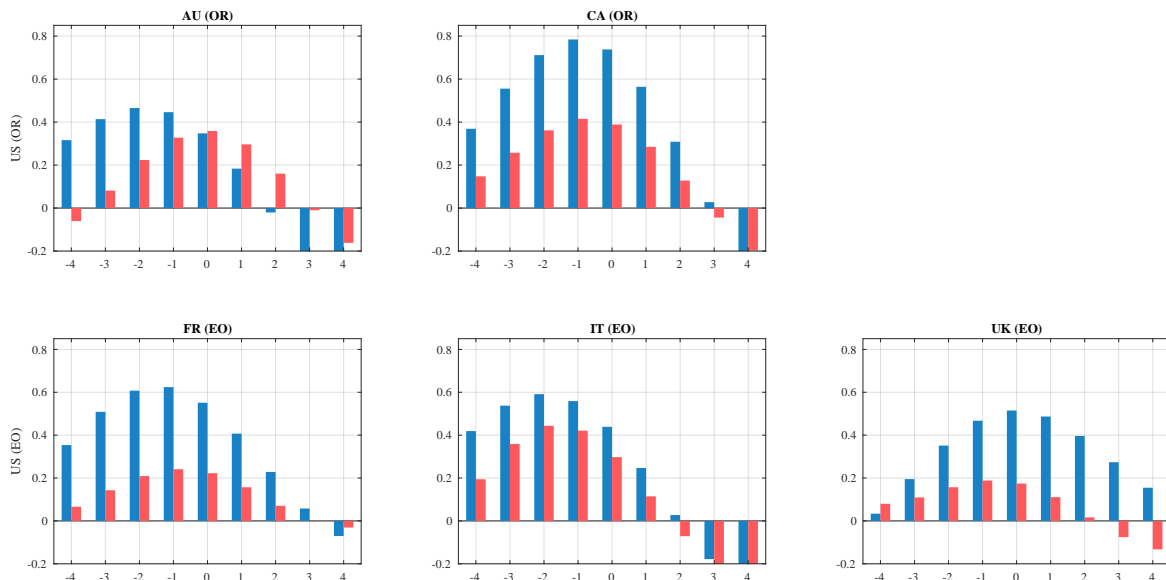
An observation that spans across all pairs is that the counterfactual series are less contemporaneously correlated than their actual counterparts. Given that the MBC shock explains the majority, but not all of domestic business-cycle fluctuations in output, it is not surprising that it also does not explain 100% of international comovement in output.¹⁶ Yet, for European country pairs, the MBC shock explains a significant fraction (about 70%) of output synchronization. If we look at the transatlantic pairs, the share of explained correlation drops to about

¹⁵See [Australian Trade and Investment Commission \(Austrade\)](#).

¹⁶The exact contributions of the output shock to the business-cycle volatility of output are displayed in Table 2.2. For CA, the UK, and the US, they are around 80%. For FR and IT, about 70%.

40% but stays elevated for Italy. We conclude that the MBC shocks in European countries are strongly linked with other European countries. If we look at the US-CA pair, we find that the MBC shock explains almost all of the comovement in the OECD EO dataset, which starts in 1981:Q1. If we look at the longer OECD OR dataset, which begins in 1961:Q2, the correlation drops, but the MBC shock still explains more than half of business-cycle synchronization. We conclude that the MBC shocks between the US and CA are also tightly linked. Our results suggest that the MBC shock, a non-inflationary aggregate demand shock, plays a significant role in international business-cycle comovement. This finding is in line with Levchenko and Pandalai-Nayar (2020), who identify transitory demand shocks that are orthogonal to technology and find that they also explain the majority of US-Canadian business-cycle comovement and Miyamoto et al. (2022), who find that the correlation between the extracted MBC shocks across the G7 countries is substantial. Our results are moreover in line with Huo et al. (2020), who suggests that non-technology shocks that move factor utilization are more promising in explaining the international business-cycle comovement than the conventional approach that relies primarily on technology-based shocks. We now investigate the share of global output comovement that originates from the US MBC shock. In Figure 2.6, we provide the dynamic correlation of counterfactual output at leads and lags between the US and the remaining countries. Large correlations on the left-hand side of the ‘0’ column indicate that fluctuations in the US are a leading indicator.

Figure 2.6: Correlation of actual and counterfactual output with the corresponding US series (VECM1)



Blue bars indicate the correlation of the bandpass-filtered series for output with US output at various leads and lags. Red bars indicate the correlation of bandpass-filtered series for counterfactual output (output as generated by the output shock) with US counterfactual output at various leads and lags.

The blue bars show that US business-cycle fluctuations are a leading factor for business-cycles in the remaining countries. The red bars indicate that the US MBC shock retains this leading character, especially for Australia, Canada, and Italy. It is, however, not the dominant leading force in France and the UK, where, as we saw above, the MBC shocks are more influenced by European dynamics.

2.4.7 Lessons from the Cross-Country Analysis

We draw three lessons from the above cross-country evidence. First, extending ACD's findings for the US, we find support for the existence of a single dominant force in each of the studied countries that explains the bulk of macroeconomic fluctuations, triggers strong comovement in the main macroeconomic variables, and is disconnected from inflation and the supply forces that drive economic activity in the long-run. Echoing Lucas (1977), we find that business-cycles are indeed all alike along the international dimension and driven by the same type of shocks; it is hence possible to study them using a unifying business-cycle theory. Our paper provides an international empirical template. We discuss how the domestic properties of the identified MBC shocks map into the existing predominantly closed economy literature in Section 2.A in the Appendix in more detail. In short, the MBC shock relates to the recent literature on non-inflationary aggregate demand shocks. A prominent example thereof is Beaudry and Portier (2014), who illustrate how non-inflationary demand-driven business-cycles may arise due to changes in perception about the future. Another example is Angeletos et al. (2018), who introduce non-inflationary demand-driven business-cycles by introducing aggregate variation in higher-order beliefs.

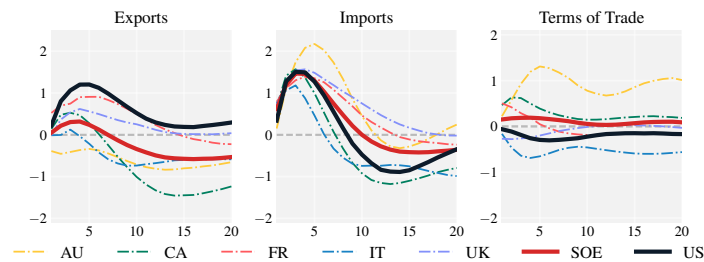
Second, we establish the existence of a MLR shock, which drives economic activity in the long but remains disconnected from economic activity in the short-run.

Third, the MBC shock, a non-inflationary aggregate demand shock, can explain a considerable fraction of international output comovement. These findings support the recent international business-cycle literature that emphasizes the importance of non-technology shocks to explain international business-cycle comovement, including Levchenko and Pandalai-Nayar (2020) and Huo et al. (2020). Our findings of international output synchronization via the main business-cycle driver prompt us to include international trade variables in our baseline model to investigate the relationship between the MBC shock and fluctuations in international trade in the next section.

2.5 Opening Up

To investigate the relationship between the drivers of fluctuations in aggregate activity and international trade, we amend the baseline domestic dataset with the terms of trade, exports, and imports.¹⁷ Including these variables does not alter the picture of the MBC shock as described in the previous section. Figure 2.7 focuses on the reaction of the newly added variables exports, imports, and the terms of trade to the MBC shock. The variance contributions of the MBC shocks are summarized in Table 2.11 for the SOE and the US.¹⁸

Figure 2.7: IRFs to the MBC shocks around the world (VECM1)



IRFs of open economy variables to the MBC shock (identified as the shock that targets unemployment).
Horizontal axis: Time horizon in quarters. Vertical axis: Percent.

Table 2.11: Variance contribution of the MBC shock around the world (VECM1)

	u	Y	h	I	C	LP	EXP	IMP	TOT	π	R
MBC SOE	69.7	50.3	49.4	51.8	36.7	25.5	17.6	36.9	11.6	11.7	28.1
MBC US	65.2	45.7	59.8	48.8	35.7	24.0	22.5	37.1	6.2	15.4	28.8

Columns give the contributions of the constructed MBC to the business-cycle volatility of the variables.

Overall, we see that the MBC shock elicits a large pro-cyclical reaction in imports peaking after about one year. Moreover, we observe a smaller procyclical reaction in exports, and hence a decrease in net exports, as exports increase by less than imports. For Australia, Canada, and Italy, we even find that after the MBC shock, their exports do not significantly increase and are even slightly negative. We observe a heterogenous, small to moderate reaction in the terms of trade. For Canada and France, there is a very short-lived increase. For Italy, the UK, and the US, the terms of trade react negatively but only barely for the latter two, and more significantly for the former. Finally, relative to the other countries, Australia's terms of trade increase quite substantially and more persistently but are also subject to more uncertainty

¹⁷We also estimated systems including net exports in percent of GDP instead of exports and imports. Because the results are quantitatively and qualitatively very similar, we omit these results for reasons of space.

¹⁸The full set of IRFs to the MBC are provided in Figure 2.16, and the variance contributions are provided in Table 2.26, and Table 2.27 in Appendix 2.C.3.

around the estimate, likely due to the fact that Australia exhibits the most volatile terms of trade in the sample.

With regard to variance contributions, the MBC shock explains a sizeable fraction of the business-cycle volatility of imports (about 40%). In contrast, the MBC shock explains only about half as much of the business-cycle volatility of exports (about 20%). Regarding the contribution to the volatility of the terms of trade, none of the facets of the MBC shock explains a noteworthy proportion. Across the board, their variance contributions at business-cycle frequencies remain low and fluctuate in the ballpark of 12%. The MBC shock hence seems largely disconnected from the nominal side of trade, which is here represented by the terms of trade.

We now turn to the properties of the shocks that target imports, exports, and the terms of trade at business-cycle frequencies. Table 2.12 displays the contemporaneous correlation between shocks. Table 2.13 provides the variance contributions and Figure 2.8 displays the IRFs to the MBC shocks, as well as the import, export, and terms of trade shocks.

We find that the import shock is strongly related to the MBC shock and likely constitutes another facet of the MBC shock. First, the import shock triggers the familiar, positive, and transitory comovement in the main macro aggregates. The only exceptions are Italy and the US, where the positive effect of the import shock is only very short-lived. Second, the import shock explains the bulk of business-cycle fluctuations of imports and about a third to half of the business-cycle volatility of the key macroeconomic variables. Third, the import shock is highly contemporaneously correlated to the shocks that constitute the MBC shock, as displayed below.

Table 2.12: Correlation of shocks (VECM1)

	u_t	Y_t	EXP_t	IMP_t	TOT_t
SOE					
EXP_t	0.28	0.52	1.00	0.38	0.27
IMP_t	0.69	0.69	0.38	1.00	0.23
TOT_t	0.20	0.28	0.27	0.23	1.00
US					
EXP_t	0.45	0.33	1.00	0.60	0.50
IMP_t	0.68	0.47	0.60	1.00	0.41
TOT_t	0.12	-0.15	0.50	0.41	1.00

While the export shock shares similar dynamics as the MBC shock in the majority of studied countries, the export shock is less strongly connected to the MBC shock than the import shock. First, the export shock does not explain a significant fraction of the key macroeconomic variables other than output and hence has a more narrow footprint than the other facets of the

MBC shock, including the import shock. Second, the export shock is less contemporaneously correlated to the shocks that constitute the MBC shock. This suggests that the MBC shock is of domestic origin and not primarily triggered by an export shock.

In line with the observations that the shocks that target the main macroeconomic aggregates explain little of the business-cycle volatility of the terms of trade, the other direction also holds true. The shock that targets the terms of trade at business-cycle frequencies explains only very little of the variation in the key real macroeconomic variables. The only noteworthy contribution of the terms of trade shock is observed for inflation in Canada (about 60%) and Australia (about 40%). In terms of IRFs, we find that for the open economies, the terms of trade shocks do share similar dynamics as the MBC shock, but the effects on unemployment and output are generally quite weak, despite sizeable effects on the terms of trade themselves. Similar to our observed disconnect between the MBC shock and the terms of trade, Miyamoto et al. (2022) observe a disconnect between the MBC shock and the real exchange rate.

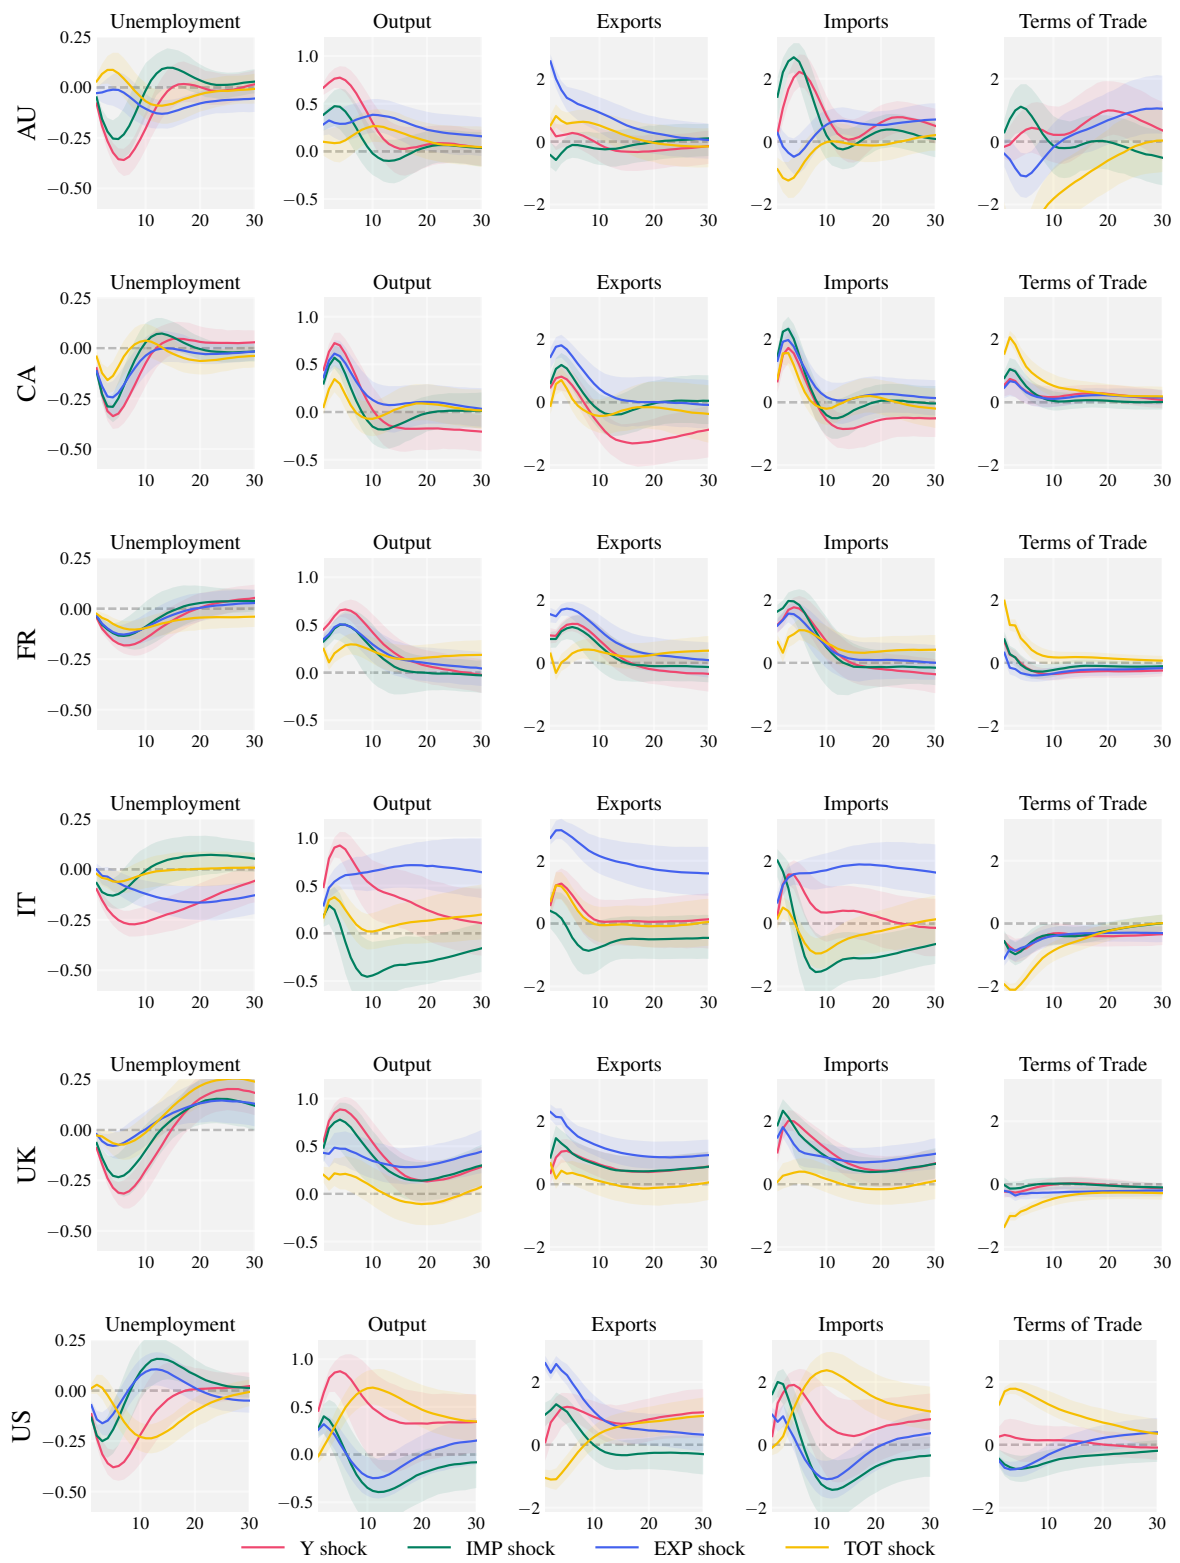
Finally, it is noteworthy how the export, import, and terms of trade shocks explain only a negligible fraction of one another and are - for the open economies - only weakly correlated.

Table 2.13: The various facets of the MBC shock, variance contributions (VECM1)

	u	Y	EXP	IMP	TOT		u	Y	EXP	IMP	TOT
AU						IT					
Unemployment	72.2	61.1	9.7	42.8	12.8	Unemployment	73.2	28.3	6.0	20.5	6.4
Output	57.0	77.0	11.1	38.8	9.6	Output	35.6	61.9	21.4	23.4	14.7
Exports	6.0	10.4	63.4	6.7	11.0	Exports	4.7	17.3	76.4	13.3	15.9
Imports	46.8	44.0	9.6	63.0	16.4	Imports	23.9	29.7	14.4	54.6	16.2
Terms of Trade	11.0	7.9	13.3	12.8	61.7	Terms of Trade	5.4	13.7	17.8	12.6	81.6
CA						UK					
Unemployment	66.4	56.2	33.5	51.3	13.8	Unemployment	73.3	63.3	13.5	39.4	13.4
Output	55.9	66.7	38.6	49.2	14.3	Output	58.2	77.4	24.7	52.5	11.4
Exports	28.8	38.6	56.4	42.6	12.3	Exports	8.5	19.0	73.4	26.7	7.2
Imports	47.9	47.0	37.3	68.2	27.8	Imports	39.9	58.8	35.0	60.0	6.6
Terms of Trade	14.1	14.6	14.5	24.4	79.1	Terms of Trade	11.0	9.9	7.1	5.9	84.1
FR						US					
Unemployment	63.4	42.7	25.2	30.4	11.4	Unemployment	65.2	45.7	22.5	37.1	6.2
Output	43.2	63.9	39.2	43.3	14.2	Output	52.6	58.6	19.9	36.8	8.9
Exports	23.4	38.1	61.0	32.5	8.7	Exports	20.2	20.8	76.7	23.3	21.1
Imports	33.0	46.4	36.8	58.1	16.7	Imports	41.8	36.4	30.4	58.9	16.2
Terms of Trade	10.8	11.0	5.4	12.3	64.6	Terms of Trade	16.1	21.3	22.7	24.8	73.7

The rows correspond to different targets in the construction of the shock. The columns give the contributions of the constructed shock to the business-cycle volatility of the variables.

Figure 2.8: Comparison of MBC shock with import, export, and terms of trade shocks, IRFs (VECM1)



IRFs of selected variables to selected shocks. Horizontal axis: Time horizon in quarters. Vertical axis: Percent.

2.5.1 Lessons from the Open Economy Anatomy

With regard to open economy anatomy, we observe the following points. First, the MBC shock affects imports positively and more than it affects exports, causing net exports to decrease.

Second, import shocks are relatively closely connected to the key real macroeconomic variables and are - in terms of variance contributions, IRFs, and correlation of shocks - similar to the MBC shock. In this respect, the import shock constitutes another facet of the MBC shock.

Third, the export shock shares similar dynamics as the MBC shock and the import shock for most countries. Apart from output, the export shock does, however, not explain a large fraction of business-cycle fluctuations of the key macroeconomic variables. In addition, the export shock is less contemporaneously correlated to the shocks that constitute the MBC shock. These observations suggest a domestic origin of the MBC shock. Fourth, similar to the disconnect between the MBC shock and inflation, we observe a disconnect between the MBC shock and the terms of trade and between the terms of trade shock and the key macroeconomic variables. If at all, the only relevant connection of the terms of trade shock seems limited to domestic inflation.

2.6 Conclusion

This paper set out to address the question: Are business cycles all alike and driven by the same type of main force in the cross-section of advanced economies? Supporting Lucas (1977), our empirical evidence shows that business cycles in the studied advanced economies are indeed all alike and driven by the same type of shock, the MBC shock. We provide an international empirical template. Within each country, the MBC shock explains the bulk of macroeconomic fluctuations and triggers strong comovement in the key macroeconomic variables. Unemployment decreases, output, investment, consumption, and hours worked increase, and net exports react countercyclically as imports increase more than exports. Moreover, the MBC shocks are disconnected from inflation and supply forces that drive economic activity in the long-run. In terms of international synchronization, the MBC shocks explain a considerable fraction of international output comovement, in line with recent research in the international business-cycle literature (Levchenko and Pandalai-Nayar (2020) and Huo et al. (2020)) that propose demand-side shocks to be important drivers of international business-cycle comovement. Concerning the link between the main driver of fluctuations in aggregate activity and international trade, the results for the six countries considered show that the main business-cycle driver is strongly connected to imports and that the import shocks are strongly connected to the main driver of fluctuations in aggregate activity. The export shock shares similar dynamics as the MBC shock and the import shock. The export shock does, however, not explain a significant fraction

of business-cycle fluctuations of the key macroeconomic variables other than output. Moreover, compared to the import shock, the export shock is less contemporaneously correlated to the other shocks that constitute the MBC shock, suggesting a domestic origin of the MBC shock. Similar to the disconnect between the MBC shock and inflation, we observe a disconnect between the MBC shock and the terms of trade. In line with recent evidence that shows a disconnect between the real exchange rate and the MBC shocks in the G7 countries (Miyamoto et al. (2022)), we conclude that the main driver of fluctuations in aggregate activity and the main driver of the terms of trade are disconnected and representatives of distinct forces.

2.A Appendix I

This section briefly summarizes ACD’s discussion of how the established domestic properties of the MBC shock map into theory. First, the disconnect of the MBC shock from TFP challenges as candidates TFP shocks, TFP news shocks, or other candidates that trigger business-cycles through endogenous movements in productivity and the real business-cycle (RBC) model. See Beaudry and Portier (2006), Lorenzoni (2009), Benhabib and Farmer (1994), Bloom et al. (2018), and Bai et al. (2017). Second, the MBC shock’s disconnect from technology and its transitory nature suit a textbook aggregate demand (AD) shock. However, the non-inflationary nature of the former also precludes the latter as a suitable candidate. Two candidates remain. The first candidate is a demand shock as featured in modern New Keynesian dynamic stochastic general equilibrium (NK DSGE) models. In this class of models, nominal rigidities introduce a non-vertical Phillips Curve, which allows demand-driven fluctuations to affect output while the reaction of inflation is dampened. Demand-driven fluctuations exist inside these models by creating a gap to unobserved flexible price allocations. The problem of accommodating non-inflationary aggregate demand fluctuations in these models arises from the standard specification of the monetary policy rule. Specifically, when the monetary authority reacts to inflationary pressure by raising the policy rate, it closes the output gap. It thus completely offsets the space in which demand-driven fluctuations were supposed to operate in the first place. Due to its non-inflationary property, the MBC shock resembles a different class of demand shock. One candidate is proposed by Beaudry and Portier (2014). Another is formulated in Angeletos et al. (2018).

Beaudry and Portier (2014), illustrate non-inflationary demand-driven business-cycles due to changes in perception about the future. In particular, Beaudry and Portier (2014) show that in a New Keynesian model augmented with labor market specialization, perceptions about future productivity can enter the natural level of output. Increases in output due to changes in perception about capital productivity then do not elicit inflationary pressure as the natural output level has also changed towards a more favorable output-inflation trade-off. In this framework, the monetary authority would want to accommodate these non-inflationary aggregate demand shocks as they create a business-cycle upswing but do not elicit any inflationary pressure.

In Angeletos et al. (2018)’s theoretical approach, demand-driven business-cycles can operate directly within the flexible price core of a NK DSGE model by allowing for aggregate variation in higher-order beliefs. Higher-order beliefs are introduced by slightly departing from rational expectations *via* the introduction of heterogeneous priors. Agents receive the same signal of a fundamental (which is taken to be TFP), knowing that it is unbiased. These agents, however, believe that the signals of *others* are biased by an exogenous variable. Shocks to this

exogenous variable then constitute confidence shocks about the short-run economic outlook of the economy. Following a positive confidence shock, firms and households have optimistic expectations of profitability and income in the short-run, which ultimately triggers a transitory macroeconomic boom in the main macroeconomic aggregates without a concomitant increase in the general level of technology. As shown in Angeletos et al. (2018), their confidence shock can replicate the comovement of the key macroeconomic variables as observed with the MBC shock. In their NK model, it moreover constitutes the dominant driver of the business-cycle and is also disconnected from inflation, TFP, and the long-run.

2.B Appendix II

This section provides further details about the data used, the trade openness index of each country, a graphical description of the observed comovement of the main macroeconomic time series over business-cycle frequencies and additional descriptive statistics of the data.

2.B.1 OECD EO

The following series are used for the OECD EO dataset. The codes refer to the ones for the OECD/EO dataset on `dbnomics`. The data are downloaded from <https://db.nomics.world/>.

Table 2.14: OECD EO dataset

Series Name	Code
Gross domestic product (volume)	GDPV.Q
Gross domestic product (deflator)	PGDP.Q
Private final consumption expenditure (volume)	CPV.Q
Gross capital formation (volume)	ITISKV.Q
Imports of goods and services (national accounts basis) (volume)	MGSV.Q
Exports of goods and services (national accounts basis) (volume)	XGSV.Q
Terms of trade, goods and services	TTRADE.Q
Labor productivity, total economy	PDTY.Q
Hours worked per worker, total economy	HRS.Q
Unemployment Rate	UNR.Q
Short-term interest rate	IRS.Q
Long-term interest rate on government bonds	IRL.Q
Total Employment (labor force survey basis)	ET.Q
Working-age population, age 15-74 (linearly interpolated from annual to quarter)	POP1574.A

2.B.2 OECD OR

The following series are used for the OECD OR dataset and are from Ohanian and Raffo (2012). The data is downloaded from <http://andrearaffo.com/araffo/Research.html>.

Table 2.15: Ohanian and Raffo dataset

Series Name	Code
Hours worked per worker	Hours
Total employment	Empl
Population aged 15 to 64	Pop

2.B.3 Further Details

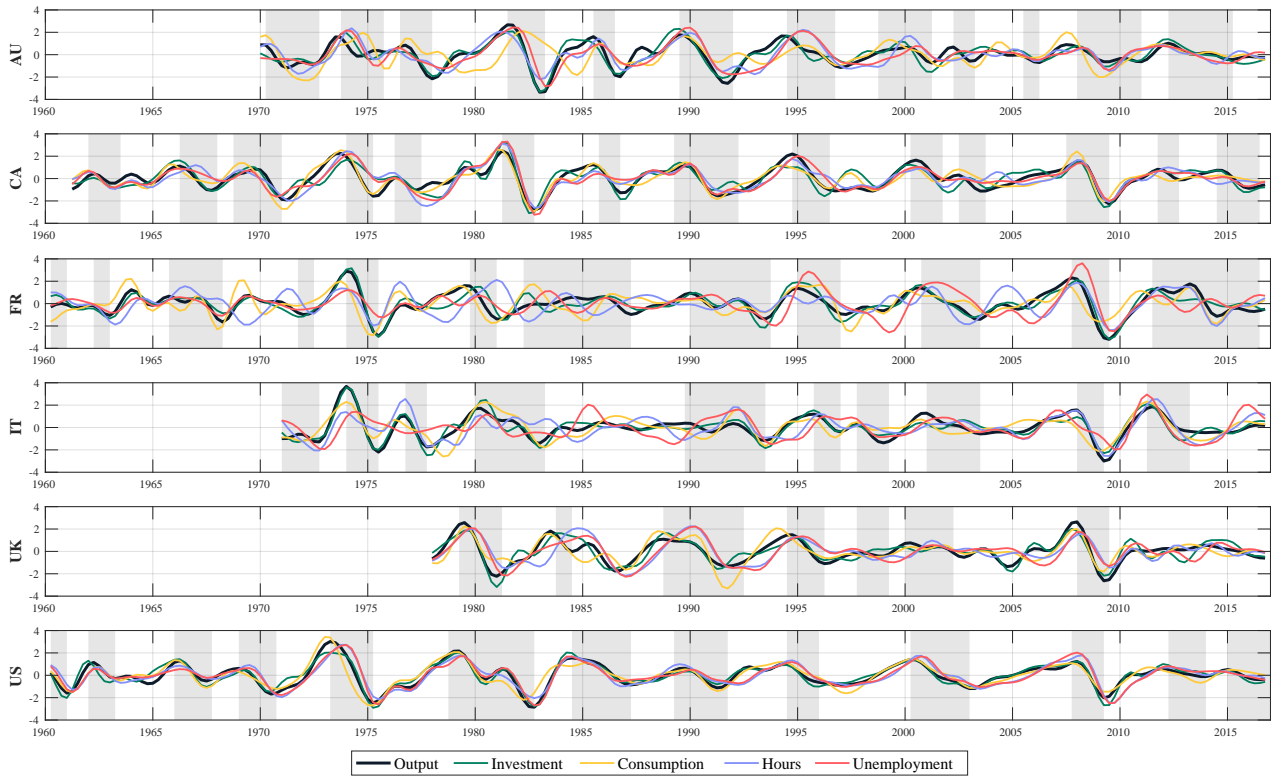
Due to data availability issues, for Canada, instead of `HRS.Q`, we use ‘Average hours worked - Total Economy’ from the `STATCAN` dataset, Code: `STATCAN/36100207/1.3.19`. For France, instead of `IRS.Q`, we use the Overnight interbank rate from the `OECD KEI` dataset. Code: `IRSTCI01.FRA.ST.Q`. Total hours per capita are calculated as: $H = h * E/P$, where h is hours worked per worker, E is total employment and P is population. The dataset referred to as `OECD EO` uniquely uses the data from `OECD EO`. The dataset referred to as `OECD OR` is the same but uses the data from Ohanian and Raffo (2012) for h , E and P . Net exports are calculated as $\frac{EXP-IMP}{Y}$, where EXP (Code: `XGS.Q`) are nominal exports of goods and services, IMP are nominal imports of goods and services (Code: `MGS.Q`), and Y is the nominal gross domestic product (Code: `GDP.Q`) from the `OECD EO` dataset. For data on capacity utilization, we use data from various sources. For Australia, the ACCI-Westpac Survey of Industrial Trends reports a capacity utilization in manufacturing, but could not be obtained for this paper. For Canada, we merge the series for industrial capacity utilization rates, by Standard Industrial Classification 1980 (SIC), from Manufacturing Industries [E], Code: `CANSIM 028-0001 (1962:25-1987:25)`, with Industrial capacity utilization rates, by industry from Manufacturing [31-33], Code: `16-10-0109-01 (1987:50-2019:00)`, from Statistics Canada (<https://www.statcan.gc.ca>). For France, Italy, UK and the US, we use capacity utilization data from business tendency surveys in manufacturing from the `OECD MEI` dataset, Code: `BSCURT02.STSA.Q`.

2.B.4 Comovement

The below figure provides the bandpass-filtered series for output, investment, consumption, hours, and unemployment for each country. The filter is parameterized to isolate the business-cycle components between 6 and 32 quarters. The resulting series are moreover normalized to

unit-variance in order to focus on comovement. The grey areas denote recessions. The data displayed are from the OECD OR dataset in order to show longer sample periods for Australia and Canada.¹⁹

Figure 2.9: Comovement in key macroeconomic series around the world



The data are detrended using the Christiano Fitzgerald approximate Bandpass-filter to isolated the business-cycle components and normalized to unit-variance to focus on comovement. Grey areas denote recessions and are obtained from the OECD based recession indicators for each from the peak through the trough for each country.

¹⁹The figure as constructed with the OECD EO dataset is very similar.

2.B.5 Moments

Below we present business-cycle moments for all variables used in our analysis.

Table 2.16: Correlations and standard deviations (bandpass-filtered, 6-32 quarters)

$\rho(\cdot, Y_t)$	AU	CA	FR	IT	UK	US	σ	AU	CA	FR	IT	UK	US
Y_t	1.00	1.00	1.00	1.00	1.00	1.00	Y_t	1.01	1.30	0.84	1.34	1.06	1.36
I_t	0.89	0.85	0.89	0.91	0.85	0.94	I_t	5.46	4.04	3.87	3.69	4.93	3.80
C_t	0.20	0.84	0.56	0.73	0.78	0.89	C_t	0.87	0.63	0.80	0.77	0.97	0.82
h_t	0.63	0.84	0.27	0.53	0.54	0.87	h_t	1.19	0.78	0.91	0.51	0.98	0.87
u_t	-0.71	-0.88	-0.61	-0.47	-0.63	-0.88	u_t	0.60	0.46	0.37	0.26	0.40	0.49
R_t	0.34	0.66	0.41	0.40	0.18	0.34	R_t	0.37	0.22	0.37	0.29	0.27	0.24
π_t	0.16	0.17	0.10	0.55	0.19	0.21	π_t	0.50	0.33	0.38	0.40	0.39	0.16
LP_t	0.54	0.83	0.87	0.89	0.80	0.63	LP_t	0.96	0.63	0.71	0.87	0.92	0.59
NX_t	-0.44	0.19	-0.19	-0.58	-0.51	-0.62	NX_t	0.94	0.65	0.68	0.67	0.54	0.27
EXP_t	0.05	0.83	0.81	0.62	0.44	0.45	EXP_t	2.84	2.59	3.24	3.11	2.56	2.44
IMP_t	0.61	0.81	0.79	0.84	0.67	0.86	IMP_t	5.03	3.36	3.74	3.16	3.00	3.28
TOT_t	0.12	0.36	-0.08	-0.59	-0.27	0.12	TOT_t	5.06	1.85	2.47	2.07	1.07	1.65

Contemporaneous correlation with output.

In absolute terms for output and relative to output for remaining variables.

2.B.6 Trade Openness Index

The trade openness index indicates how much a country is exposed to international trade. In Table 2.17 we can see the degree of trade openness among the six countries considered. In all three subsamples considered, Canada ranks as the most open country, followed by France, the UK, and Italy. Australia consistently ranks fifth, and the US ranks last with the lowest trade openness of only about 25%. Table 2.18 and Table 2.19 display the relative size of exports and imports to GDP and exhibit a very similar pattern.

Table 2.17: Trade Openness Index: Trade in % of GDP

Average 1971-2019		Average 1981-2019		2019	
1. Canada	(59%)	1. Canada	(63%)	1. Canada	(65%)
2. United Kingdom	(53%)	2. United Kingdom	(53%)	2. France	(65%)
3. France	(48%)	3. France	(51%)	3. United Kingdom	(63%)
4. Italy	(45%)	4. Italy	(47%)	4. Italy	(60%)
5. Australia	(36%)	5. Australia	(38%)	5. Australia	(46%)
6. United States	(22%)	6. United States	(23%)	6. United States	(26%)

Numbers in brackets are trade in % of GDP, calculated as the sum of exports and imports of goods and services measured as a share of gross domestic product. Source: World Bank, indicator code `NE.TRD.GNFS.ZS`.

Table 2.18: Exports in % of GDP

Average 1971-2019		Average 1981-2019		2019	
1. Canada	(30%)	1. Canada	(32%)	1. Canada	(32%)
2. United Kingdom	(26%)	2. United Kingdom	(26%)	2. France	(32%)
3. France	(24%)	3. France	(26%)	3. Italy	(32%)
4. Italy	(23%)	4. Italy	(24%)	4. United Kingdom	(31%)
5. Australia	(18%)	5. Australia	(18%)	5. Australia	(24%)
6. United States	(10%)	6. United States	(10%)	6. United States	(12%)

Numbers in brackets are exports in % of GDP. Exports of goods and services represent the value of all goods and other market services provided to the rest of the world. Source: World Bank, indicator code `NE.EXP.GNFS.ZS`.

Table 2.19: Imports in % of GDP

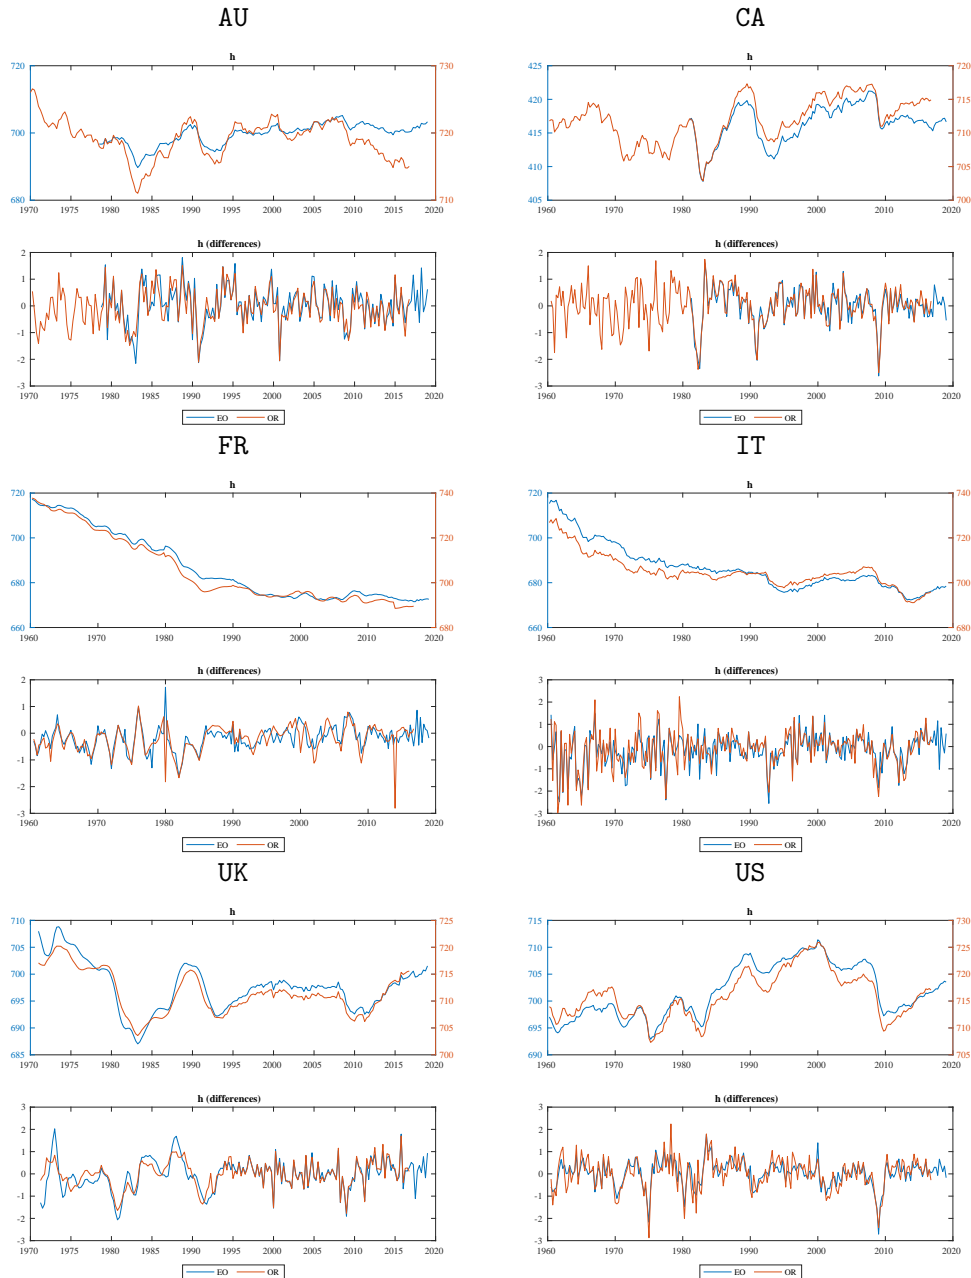
Average 1971-2019		Average 1981-2019		2019	
1. Canada	(29%)	1. Canada	(31%)	1. Canada	(34%)
2. United Kingdom	(27%)	2. United Kingdom	(27%)	2. France	(33%)
3. France	(24%)	3. France	(26%)	3. United Kingdom	(32%)
4. Italy	(22%)	4. Italy	(23%)	4. Italy	(28%)
5. Australia	(18%)	5. Australia	(20%)	5. Australia	(22%)
6. United States	(12%)	6. United States	(13%)	6. United States	(15%)

Numbers in brackets are imports in % of GDP. Imports of goods and services represent the value of all goods and other market services received from the rest of the world. Source: World Bank, indicator code `NE.IMP.GNFS.ZS`.

2.B.7 Comparison OECD EO and OR

This section provides a comparison of the measure for hours worked between the OECD EO and OECD OR datasets. As noted in Section 2.3.1 in the main text, the OECD OR dataset allows us to significantly extend sample lengths for Australia and Canada. For overlapping periods, the resulting series from each dataset are near-identical.

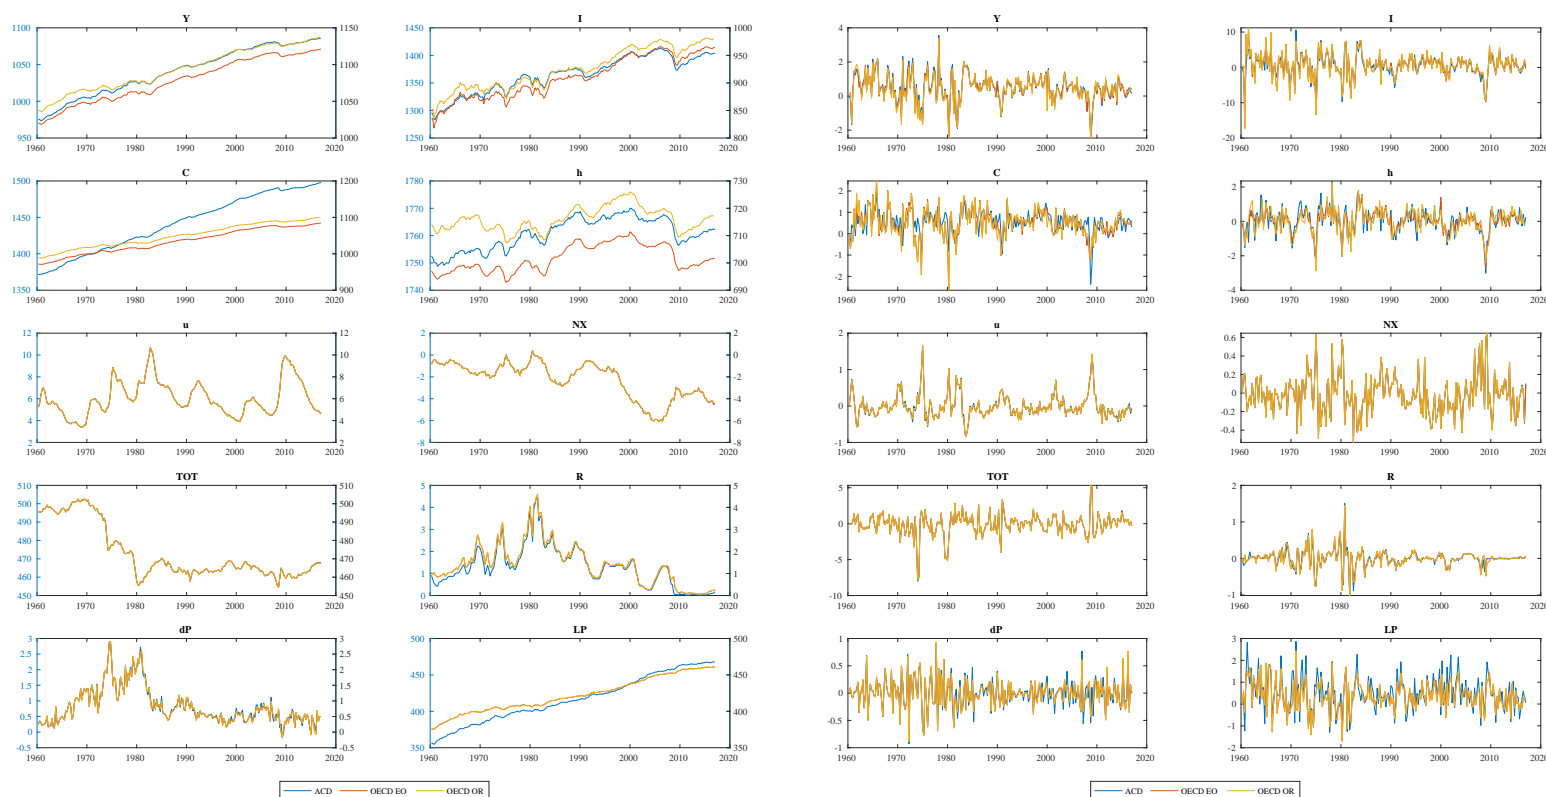
Figure 2.10: Comparison of hours



2.B.8 Comparison US: Comparison EO, OR and ACD

In this section, we provide for the US a comparison between the datasets stemming from the OECD EO dataset, the OECD OR dataset and the ACD dataset as used in Angeletos et al. (2020). The resulting series from each dataset are near-identical. The left panel plots the data in (log) levels, except for inflation. The right panel plots differences.

Figure 2.11: Comparison of datasets for the US.



2.C Appendix III

This section provides additional results as indicated in the main text.

2.C.1 Cross-Country Analysis: VECM1

Table 2.20: Correlations of conditional times series (VECM1)

Variable	Shock			
	I	C	h	u
AU				
Unemployment	0.920	0.347	0.880	0.898
Output	0.973	0.866	0.970	0.962
Investment	0.907	0.649	0.929	0.929
Consumption	0.896	0.098	0.784	0.904
Hours Worked	0.932	0.242	0.861	0.889
CA				
Unemployment	0.937	0.911	0.989	0.981
Output	0.964	0.964	0.997	0.988
Investment	0.921	0.940	0.992	0.975
Consumption	0.943	0.839	0.985	0.962
Hours Worked	0.937	0.921	0.986	0.981
FR				
Unemployment	0.956	0.821	0.482	0.845
Output	0.986	0.813	0.425	0.799
Investment	0.936	0.802	0.498	0.825
Consumption	0.957	0.861	0.781	0.870
Hours Worked	0.996	0.817	0.871	0.972
IT				
Unemployment	0.904	0.654	0.676	0.678
Output	0.974	0.860	0.968	0.800
Investment	0.891	0.852	0.951	0.738
Consumption	0.934	0.687	0.919	0.715
Hours Worked	0.955	0.752	0.845	0.781
UK				
Unemployment	0.916	0.983	0.840	0.933
Output	0.990	0.991	0.970	0.988
Investment	0.887	0.985	0.881	0.939
Consumption	0.973	0.938	0.955	0.980
Hours Worked	0.871	0.989	0.758	0.889
US				
Unemployment	0.991	0.880	0.919	0.905
Output	0.995	0.937	0.976	0.959
Investment	0.983	0.917	0.969	0.954
Consumption	0.980	0.906	0.962	0.951
Hours Worked	0.987	0.862	0.924	0.893

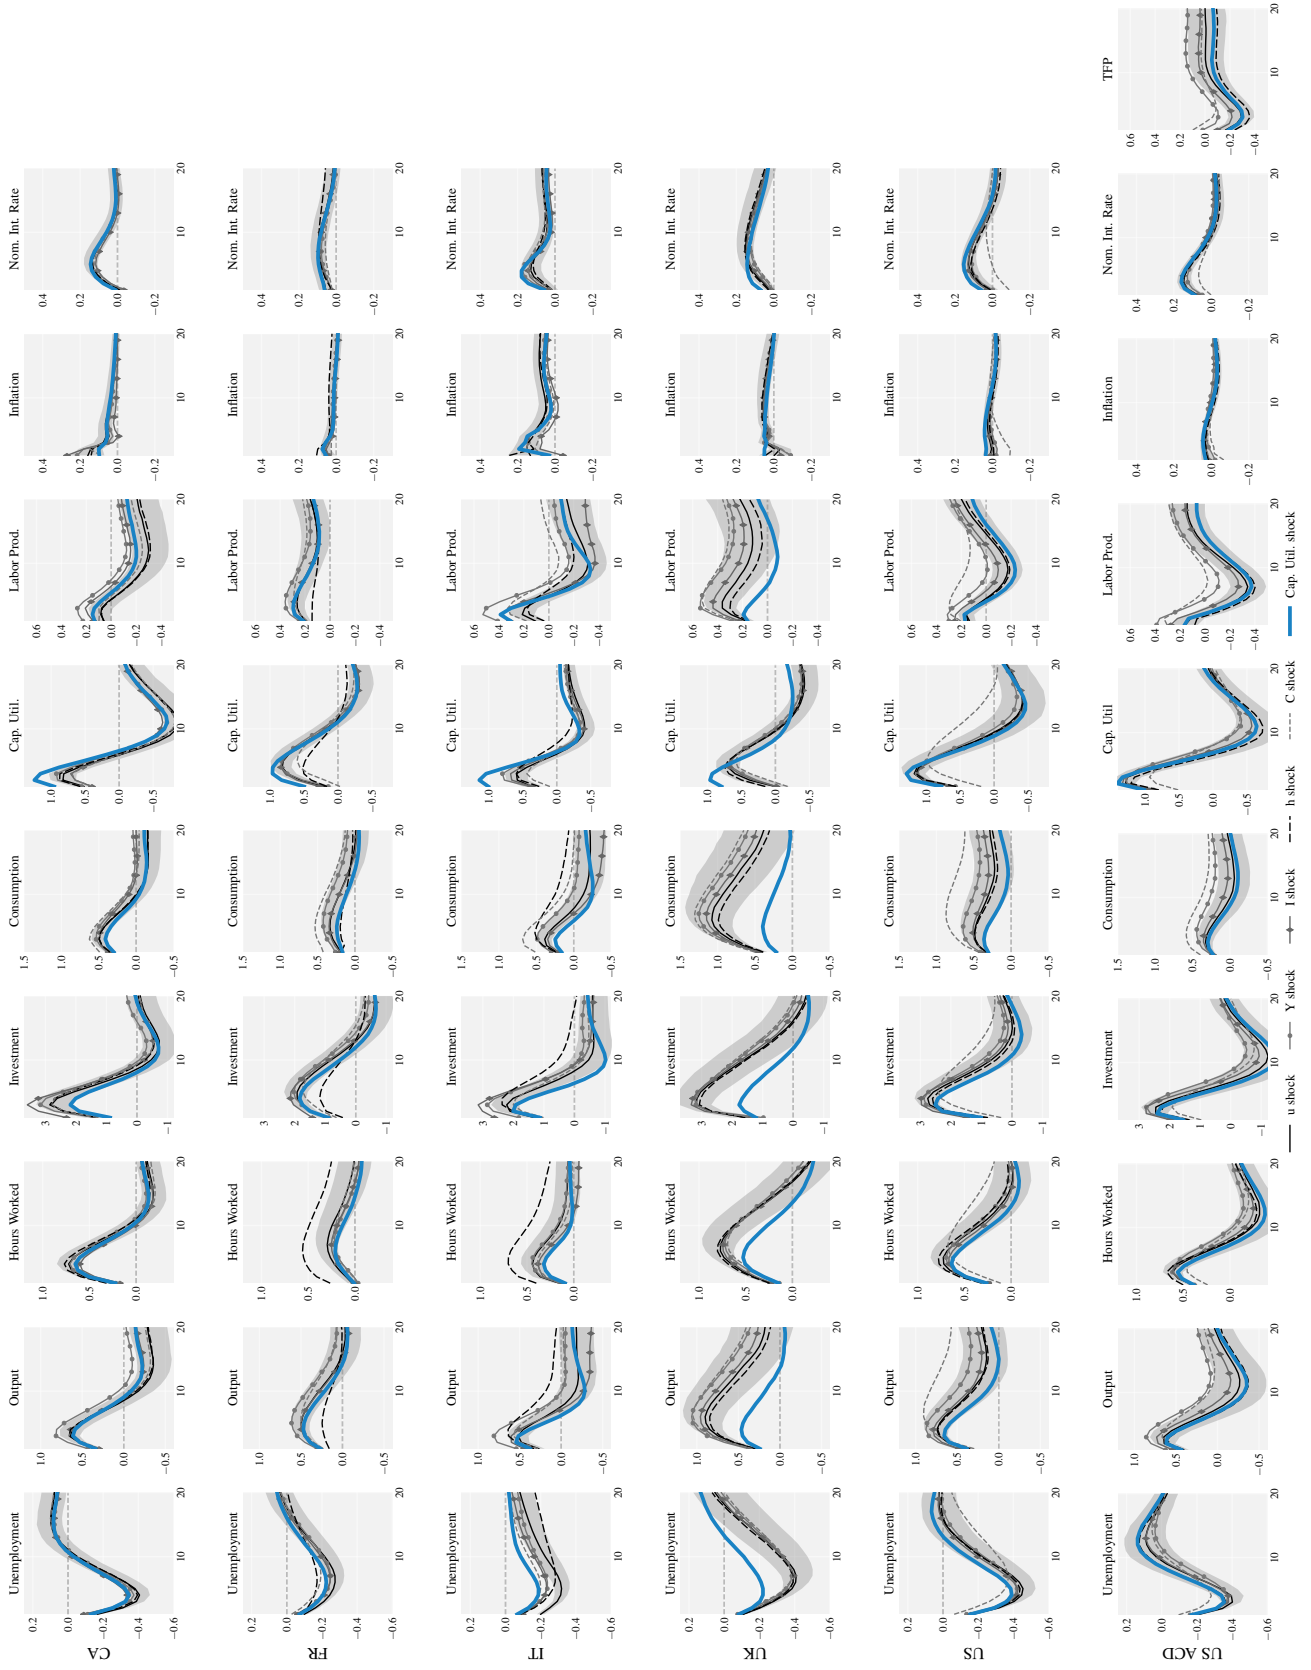
Each row reports the correlation between each bandpass-filtered variable as predicted by the output shock and that predicted by the other facets of the MBC shock.

Table 2.21: The various facets of the MBC shock, variance contributions (VECM 1)

	u	Y	h	I	C	CU	LP	π	R
CA									
Unemployment	85.0	68.0	78.9	55.7	65.3	61.5	27.1	13.6	46.9
Output	70.5	80.1	71.3	61.5	61.7	57.8	40.1	17.1	46.0
Hours Worked	79.6	69.1	84.7	55.0	63.7	62.4	27.4	11.4	43.5
Investment	58.4	61.5	57.3	80.2	53.8	39.5	28.9	21.7	35.9
Consumption	66.7	61.3	65.0	53.8	82.8	48.3	24.7	10.2	37.6
Cap. Util.	62.3	57.5	63.7	39.2	48.0	80.6	26.1	9.0	47.6
FR									
Unemployment	77.1	60.7	32.8	62.6	41.3	59.0	42.0	11.5	20.4
Output	62.3	76.3	24.8	62.8	52.6	59.7	66.8	11.3	17.4
Hours Worked	27.9	16.4	69.9	23.5	14.6	22.0	13.0	16.5	12.8
Investment	65.2	64.0	29.0	74.7	44.6	60.5	51.7	9.8	18.4
Consumption	40.7	49.3	21.2	43.1	77.2	33.2	37.9	5.7	9.5
Cap. Util.	58.5	57.9	25.8	59.1	29.5	73.4	51.6	14.3	22.3
IT									
Unemployment	74.6	32.8	39.0	38.0	28.3	33.6	20.5	8.8	13.8
Output	40.6	62.2	40.5	55.1	35.4	47.1	47.8	10.1	27.9
Hours Worked	42.1	33.6	79.8	33.0	29.5	24.6	11.6	10.1	11.3
Investment	42.9	51.3	36.3	68.5	30.5	39.5	39.6	6.0	22.0
Consumption	34.9	36.1	34.7	32.6	62.1	20.6	21.8	5.9	17.8
Cap. Util.	31.5	39.7	23.6	36.0	19.8	76.5	37.7	12.0	30.9
UK									
Unemployment	78.9	70.7	73.5	70.5	71.1	48.2	38.5	12.4	40.6
Output	66.1	85.2	61.1	68.1	81.5	41.7	61.7	16.3	29.3
Hours Worked	74.7	66.1	77.2	66.9	63.2	50.8	33.6	11.0	40.4
Investment	73.7	76.6	68.7	75.0	76.7	45.6	46.9	13.6	36.7
Consumption	64.0	79.0	56.5	65.7	88.0	36.3	54.7	15.5	27.0
Cap. Util.	43.7	36.4	47.5	34.9	23.9	65.5	24.5	11.5	43.5
US									
Unemployment	72.2	53.5	64.4	59.8	34.7	55.7	33.8	10.1	29.6
Output	60.6	64.6	56.9	61.1	46.2	51.1	37.0	10.7	23.9
Hours Worked	66.9	52.5	69.1	55.1	36.6	54.0	34.1	9.9	26.5
Investment	63.9	57.9	56.4	67.0	33.5	53.3	36.2	9.4	31.2
Consumption	41.0	52.2	42.0	40.0	68.8	32.0	26.5	26.9	11.5
Cap. Util.	61.5	50.9	57.4	55.3	29.5	64.6	42.1	16.7	38.9

The rows correspond to different targets in the construction of the shock. The columns give the contributions of the constructed shock to the business-cycle volatility of the variables.

Figure 2.12: The various facets of the MBC shock, IRFs (VECM1)



2.C.2 Cross-Country Analysis: VAR

Existence of a Main business-cycle Shock

Table 2.22: The various facets of the MBC shock, variance contributions (VAR)

	u	Y	h	I	C	LP	π	R
AU								
Unemployment	76.0	42.9	40.1	54.4	14.9	19.8	8.9	20.9
Output	37.8	87.2	27.2	43.4	30.3	53.1	5.5	9.9
Hours Worked	31.8	19.3	76.5	26.6	3.4	7.4	6.7	10.8
Investment	49.3	43.3	34.7	92.5	17.8	23.8	4.3	18.3
Consumption	4.8	9.4	2.7	2.0	67.3	6.8	1.7	5.8
CA								
Unemployment	82.4	30.6	53.6	28.2	24.7	6.6	6.7	24.3
Output	30.5	70.7	34.2	32.0	31.5	43.9	11.0	6.0
Hours Worked	56.4	31.9	80.6	31.2	21.0	6.1	5.5	13.3
Investment	30.4	29.6	32.6	81.8	20.1	11.4	12.8	14.2
Consumption	21.1	27.5	19.2	18.6	76.3	9.5	2.7	4.2
FR								
Unemployment	83.0	23.2	10.8	32.2	4.5	11.5	4.6	5.2
Output	25.1	92.5	4.2	61.0	38.6	82.9	3.7	5.4
Hours Worked	10.3	3.6	87.7	3.7	3.1	2.8	5.0	2.9
Investment	38.3	61.3	4.1	91.9	10.6	51.9	2.2	5.2
Consumption	3.9	39.1	2.9	10.3	91.7	34.3	5.0	1.0
IT								
Unemployment	76.9	12.9	10.0	16.4	8.3	4.2	3.6	4.4
Output	17.8	79.3	27.4	35.5	31.1	52.6	4.7	5.1
Hours Worked	20.2	23.2	93.3	20.1	16.1	4.3	7.3	6.1
Investment	25.8	36.7	20.6	73.3	8.0	24.4	4.0	8.0
Consumption	6.3	28.1	13.4	8.0	75.5	15.1	2.5	2.1
UK								
Unemployment	67.7	30.8	53.7	50.1	23.6	7.9	8.2	11.3
Output	24.6	79.2	21.3	29.8	49.5	57.3	14.7	3.8
Hours Worked	48.4	23.2	81.2	49.9	18.2	8.2	6.9	12.8
Investment	39.3	27.7	45.8	84.7	19.2	10.3	6.8	12.3
Consumption	15.8	48.3	10.0	16.3	79.5	37.0	15.7	2.2
US								
Unemployment	85.9	46.4	58.2	48.7	22.0	13.9	7.6	23.4
Output	47.3	83.8	40.4	63.3	42.5	40.0	1.9	12.8
Hours Worked	55.9	38.0	89.0	36.3	19.9	10.5	2.6	16.1
Investment	47.5	58.6	36.4	90.7	17.3	27.6	3.2	18.5
Consumption	21.0	46.9	21.3	19.1	77.1	20.8	1.3	1.3

The rows correspond to different targets in the construction of the shock. The columns give the contributions of the constructed shock to the business-cycle volatility of the variables.

Figure 2.13: The various facets of the MBC shock, IRFs (VAR)


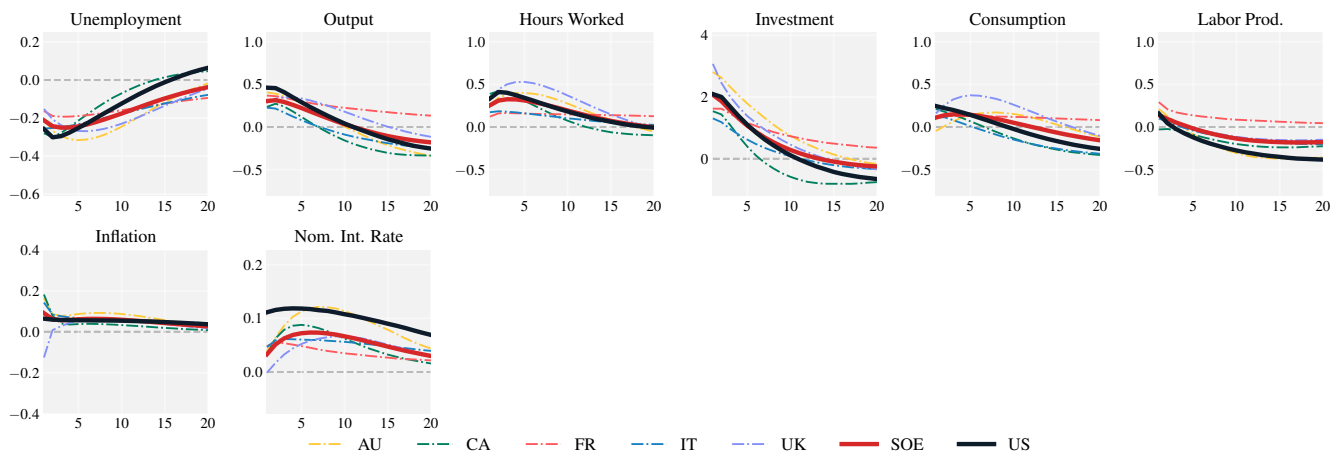


Figure 2.14: IRFs to the MBC shocks around the world (VAR)



IRFs of all the variables to the MBC shock (identified as the shock that targets unemployment). Horizontal axis: Time horizon in quarters. Vertical axis: Percent.

Table 2.23: Average correlation of conditional time series (VAR)

	AU	CA	FR	IT	UK	US
Unemployment	0.66	0.79	0.54	0.68	0.71	0.95
Output	0.98	0.97	0.98	0.99	0.96	1.00
Investment	0.84	0.91	0.86	0.90	0.84	0.98
Consumption	0.61	0.83	0.78	0.80	0.83	0.95
Hours Worked	0.62	0.83	0.67	0.86	0.73	0.94

Existence of a Main Long-Run Shock

Table 2.24: Long-run shocks, contributions at long-run frequencies ($80-\infty$ Q) (VAR)

	AU	CA	FR	IT	UK	US (OR)
Output	99.6	96.8	88.4	83.8	97.8	96.5
Investment	97.8	92.1	33.6	44.4	88.1	92.3
Consumption	99.4	98.4	85.0	83.9	97.0	96.2
Labor Productivity	99.3	97.5	86.7	80.9	96.1	82.9

The rows correspond to different targets in the construction of the shock. The columns give the contributions of the constructed shock to the long-run volatility of the variables.

Figure 2.15: The various facets of the MLR shock, IRFs (VAR)




Table 2.25: Long-run shocks, contributions at business-cycle frequencies (VAR)

	u	Y	h	I	C	LP	π	R
AU								
Output	13.9	8.2	6.5	10.7	12.2	15.2	6.1	8.4
Investment	17.9	8.2	9.0	18.3	12.0	13.5	5.7	10.3
Consumption	13.2	8.6	6.6	10.4	13.5	15.8	6.0	8.1
Labor Productivity	15.6	7.8	7.2	12.1	12.9	15.1	6.3	8.6
CA								
Output	13.4	12.8	20.0	6.5	11.9	39.7	2.3	16.2
Investment	16.7	16.7	18.8	9.4	13.9	36.1	2.7	18.1
Consumption	14.9	10.6	19.7	6.0	14.0	32.6	2.3	15.6
Labor Productivity	14.2	10.0	21.4	5.5	11.7	34.1	2.3	14.0
FR								
Output	7.4	40.7	7.1	20.7	32.6	38.9	3.6	8.6
Investment	16.4	24.1	6.7	31.7	13.4	18.0	6.3	26.0
Consumption	5.4	33.0	7.1	13.6	36.7	31.4	4.3	10.8
Labor Productivity	4.5	29.7	6.9	10.9	38.4	35.2	3.5	8.8
IT								
Output	15.7	10.0	22.7	15.5	8.5	11.5	5.6	9.6
Investment	12.9	22.3	15.7	23.8	15.5	20.1	8.9	23.5
Consumption	17.3	8.6	22.4	15.3	8.4	10.1	5.0	8.6
Labor Productivity	12.1	6.2	19.3	9.9	6.9	12.2	4.3	9.0
UK								
Output	3.8	15.6	8.0	5.1	35.3	37.8	10.1	7.6
Investment	7.4	21.3	10.9	10.5	35.6	29.3	10.1	7.2
Consumption	3.9	15.6	6.8	5.0	37.4	34.8	10.0	6.6
Labor Productivity	4.1	11.2	10.7	5.3	31.9	35.3	9.2	7.1
US (OR)								
Output	8.1	13.6	11.1	8.9	13.8	9.1	10.3	10.9
Investment	10.1	14.9	12.2	10.5	15.7	11.3	15.0	14.7
Consumption	8.5	12.9	11.3	8.9	13.7	9.1	9.6	10.8
Labor Productivity	19.2	12.2	17.0	14.1	11.2	18.3	17.1	22.2

The rows correspond to different targets in the construction of the shock. The columns give the contributions of the constructed shock to the business-cycle volatility of the variables.

2.C.3 Opening Up: VECM1

Table 2.26: The various facets of the MBC shock, variance contributions (VECM 1)

	u	Y	h	I	C	LP	EXP	IMP	TOT	π	R
AU											
Unemployment	72.2	61.1	63.6	63.7	25.0	26.7	9.7	42.8	12.8	13.0	42.5
Output	57.0	77.0	46.9	57.6	34.0	39.9	11.1	38.8	9.6	10.5	31.8
Hours Worked	62.0	48.3	73.7	52.1	15.8	20.1	7.7	38.6	15.9	12.2	32.0
Investment	57.8	56.3	49.3	78.9	24.2	24.7	9.2	54.0	11.9	10.0	38.5
Consumption	16.3	22.0	11.7	13.8	59.6	14.6	5.1	23.0	7.7	8.8	17.0
Labor Prod.	18.0	44.1	12.3	19.5	26.6	67.7	14.5	13.0	16.5	18.4	18.4
Exports	6.0	10.4	4.8	4.9	6.4	12.6	63.4	6.7	11.0	9.0	10.8
Imports	46.8	44.0	43.7	65.0	27.7	17.7	9.6	63.0	16.4	12.0	35.9
Terms of Trade	11.0	7.9	12.6	7.4	6.8	16.2	13.3	12.8	61.7	41.3	10.4
Inflation	9.2	10.1	7.6	7.0	8.1	18.1	10.0	11.1	37.6	67.4	13.4
Nom. Int. Rate	51.6	40.7	39.9	49.9	26.4	27.6	15.3	39.2	13.6	18.8	57.2
CA											
Unemployment	66.4	56.2	64.1	52.9	48.8	27.0	33.5	51.3	13.8	7.7	35.8
Output	55.9	66.7	55.5	56.4	48.5	43.9	38.6	49.2	14.3	8.9	31.6
Hours Worked	62.2	54.8	70.2	49.6	50.4	25.1	31.4	47.2	9.9	6.1	30.3
Investment	49.0	52.2	47.1	70.4	40.8	30.8	27.4	56.0	27.7	14.5	33.5
Consumption	48.5	49.1	51.3	45.1	68.0	25.5	22.5	44.1	9.4	5.2	25.6
Labor Prod.	20.5	42.8	16.6	28.8	22.5	66.5	28.1	20.5	12.4	10.3	12.7
Exports	28.8	38.6	27.5	27.6	18.8	28.7	56.4	42.6	12.3	11.2	17.6
Imports	47.9	47.0	46.0	56.8	41.0	22.6	37.3	68.2	27.8	14.9	33.1
Terms of Trade	14.1	14.6	12.3	24.7	11.0	10.6	14.5	24.4	79.1	62.3	13.6
Inflation	8.2	9.5	6.9	14.1	6.4	11.2	11.7	14.5	60.7	80.9	10.4
Nom. Int. Rate	38.3	34.2	33.0	35.5	27.7	15.7	23.6	37.9	13.7	9.7	60.5
FR											
Unemployment	63.4	42.7	16.4	40.7	29.1	29.2	25.2	30.4	11.4	7.5	13.1
Output	43.2	63.9	10.7	49.0	49.1	59.5	39.2	43.3	14.2	9.5	16.7
Hours Worked	18.5	10.8	66.2	13.4	11.6	6.6	7.1	12.7	8.0	5.8	10.2
Investment	42.8	52.0	13.3	60.5	33.7	45.2	28.3	46.3	15.8	9.6	16.4
Consumption	27.0	44.6	11.1	30.8	72.6	42.4	14.7	32.0	14.0	4.6	8.3
Labor Prod.	26.4	54.6	6.4	37.6	40.0	69.2	33.1	34.2	11.6	8.0	14.3
Exports	23.4	38.1	5.8	23.5	13.8	36.7	61.0	32.5	8.7	17.3	18.5
Imports	33.0	46.4	15.0	46.7	35.9	41.4	36.8	58.1	16.7	12.9	21.8
Terms of Trade	10.8	11.0	9.6	13.6	14.3	11.2	5.4	12.3	64.6	13.4	10.1
Inflation	4.3	5.4	7.3	5.8	3.4	6.9	13.6	9.6	16.0	74.2	20.5
Nom. Int. Rate	16.8	20.7	19.0	20.5	10.9	16.9	27.1	28.6	20.2	19.4	73.5

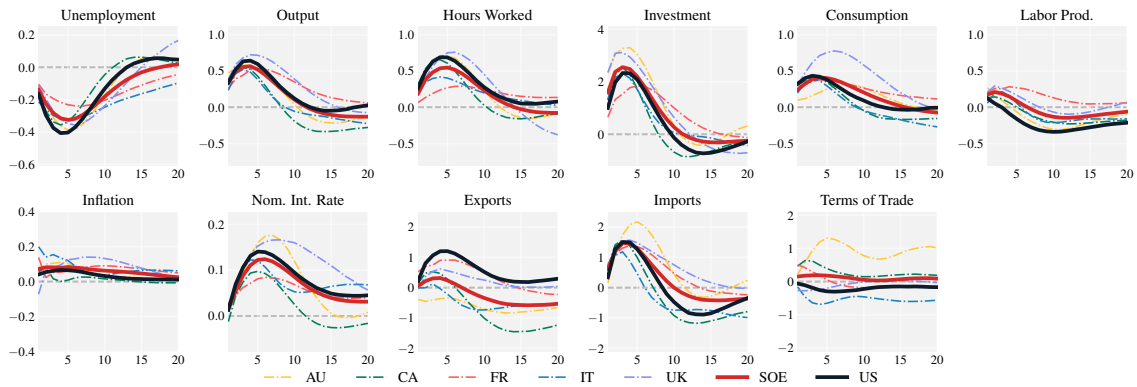
The rows correspond to different targets in the construction of the shock. The columns give the contributions of the constructed shock to the business-cycle volatility of the variables.

Table 2.27: The various facets of the MBC shock, variance contributions (VECM 1)

	u	Y	h	I	C	LP	EXP	IMP	TOT	π	R
IT											
Unemployment	73.2	28.3	34.9	37.5	24.0	13.5	6.0	20.5	6.4	11.5	11.9
Output	35.6	61.9	39.8	42.4	39.5	44.6	21.4	23.4	14.7	12.0	13.6
Hours Worked	40.7	31.5	78.2	29.8	29.9	9.1	5.6	11.7	3.7	10.2	6.7
Investment	45.1	44.4	30.4	58.8	24.3	31.6	11.4	30.7	10.0	12.0	18.0
Consumption	29.5	39.2	35.0	27.7	62.0	22.7	6.1	25.6	5.5	6.1	11.0
Labor Prod.	12.9	46.1	11.5	24.0	23.9	59.5	30.0	16.4	21.3	10.8	11.4
Exports	4.7	17.3	6.3	4.9	10.5	24.1	76.4	13.3	15.9	5.8	3.3
Imports	23.9	29.7	11.4	39.0	28.8	26.1	14.4	54.6	16.2	14.0	32.2
Terms of Trade	5.4	13.7	4.0	9.5	4.7	20.3	17.8	12.6	81.6	28.2	23.5
Inflation	9.1	8.3	10.7	7.5	5.9	9.0	5.4	8.1	18.7	70.8	33.0
Nom. Int. Rate	11.6	21.4	8.9	24.9	14.4	20.1	10.4	29.7	21.0	33.7	76.4
UK											
Unemployment	73.3	63.3	68.0	64.4	56.4	31.1	13.5	39.4	13.4	18.9	37.0
Output	58.2	77.4	52.1	61.5	67.4	52.4	24.7	52.5	11.4	19.9	27.8
Hours Worked	67.1	56.5	75.1	62.3	50.1	26.3	10.1	34.5	13.0	18.8	32.7
Investment	61.1	63.1	59.5	74.6	53.5	37.2	13.1	43.3	9.1	17.2	34.7
Consumption	50.6	67.6	44.5	52.1	76.7	44.9	12.1	43.0	7.9	14.3	20.3
Labor Prod.	27.6	57.2	20.9	36.2	48.1	68.2	20.3	37.9	10.4	19.8	11.6
Exports	8.5	19.0	7.1	8.9	7.0	15.8	73.4	26.7	7.2	6.7	11.4
Imports	39.9	58.8	35.6	49.9	47.1	38.5	35.0	60.0	6.6	13.0	25.7
Terms of Trade	11.0	9.9	10.5	6.6	7.1	10.0	7.1	5.9	84.1	24.0	14.2
Inflation	11.5	14.5	12.3	10.4	12.2	18.9	4.7	9.2	20.5	67.7	8.6
Nom. Int. Rate	33.8	23.2	25.7	31.4	13.8	14.6	14.5	21.2	10.6	19.4	65.8
US											
Unemployment	65.2	45.7	59.8	48.8	35.7	24.0	22.5	37.1	6.2	15.4	28.8
Output	52.6	58.6	52.0	53.6	51.3	29.2	19.9	36.8	8.9	8.3	20.7
Hours Worked	59.8	45.2	65.8	45.0	37.8	20.1	21.4	32.3	5.6	8.4	21.3
Investment	50.3	48.3	46.0	61.9	32.5	24.2	17.9	39.9	6.1	11.1	29.3
Consumption	40.6	51.4	42.9	41.0	63.7	24.9	12.3	34.9	14.7	8.4	11.5
Labor Prod.	16.2	26.3	14.0	18.4	30.6	60.7	8.3	21.9	18.6	21.5	12.1
Exports	20.2	20.8	20.2	17.2	14.8	16.2	76.7	23.3	21.1	15.6	20.2
Imports	41.8	36.4	36.3	42.0	29.6	24.1	30.4	58.9	16.2	26.3	31.6
Terms of Trade	16.1	21.3	14.3	15.7	24.9	20.9	22.7	24.8	73.7	33.9	13.3
Inflation	19.9	20.6	15.2	17.5	22.0	27.5	15.6	27.3	32.2	82.6	22.0
Nom. Int. Rate	28.9	28.3	26.1	30.7	19.1	23.2	22.1	28.9	12.3	13.9	67.1

The rows correspond to different targets in the construction of the shock. The columns give the contributions of the constructed shock to the business-cycle volatility of the variables.

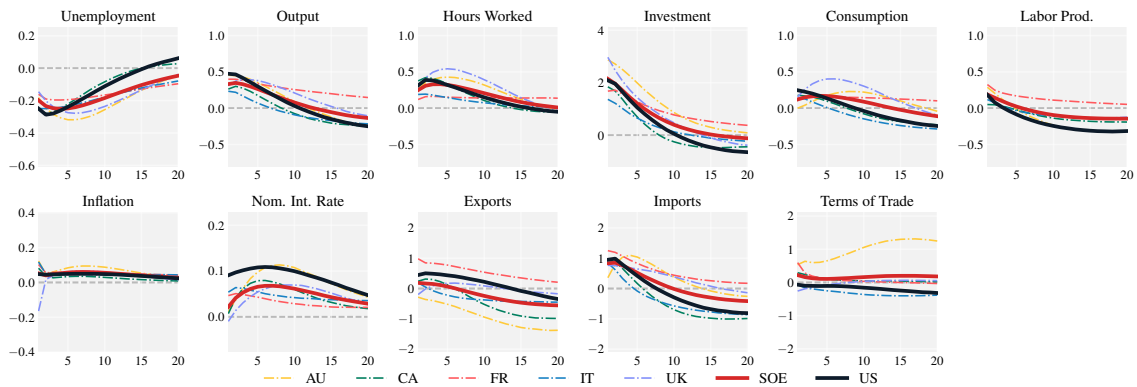
Figure 2.16: IRFs to the MBC shocks around the world (VECM1)



IRFs of all the variables to the MBC shock (identified as the shock that targets unemployment). Horizontal axis: Time horizon in quarters. Vertical axis: Percent.

2.C.4 Opening Up: VAR

Figure 2.17: IRFs to the MBC shocks around the world (VAR)



IRFs of all the variables to the MBC shock (identified as the shock that targets unemployment). Horizontal axis: Time horizon in quarters. Vertical axis: Percent.

Table 2.28: Correlation of shocks (VAR)

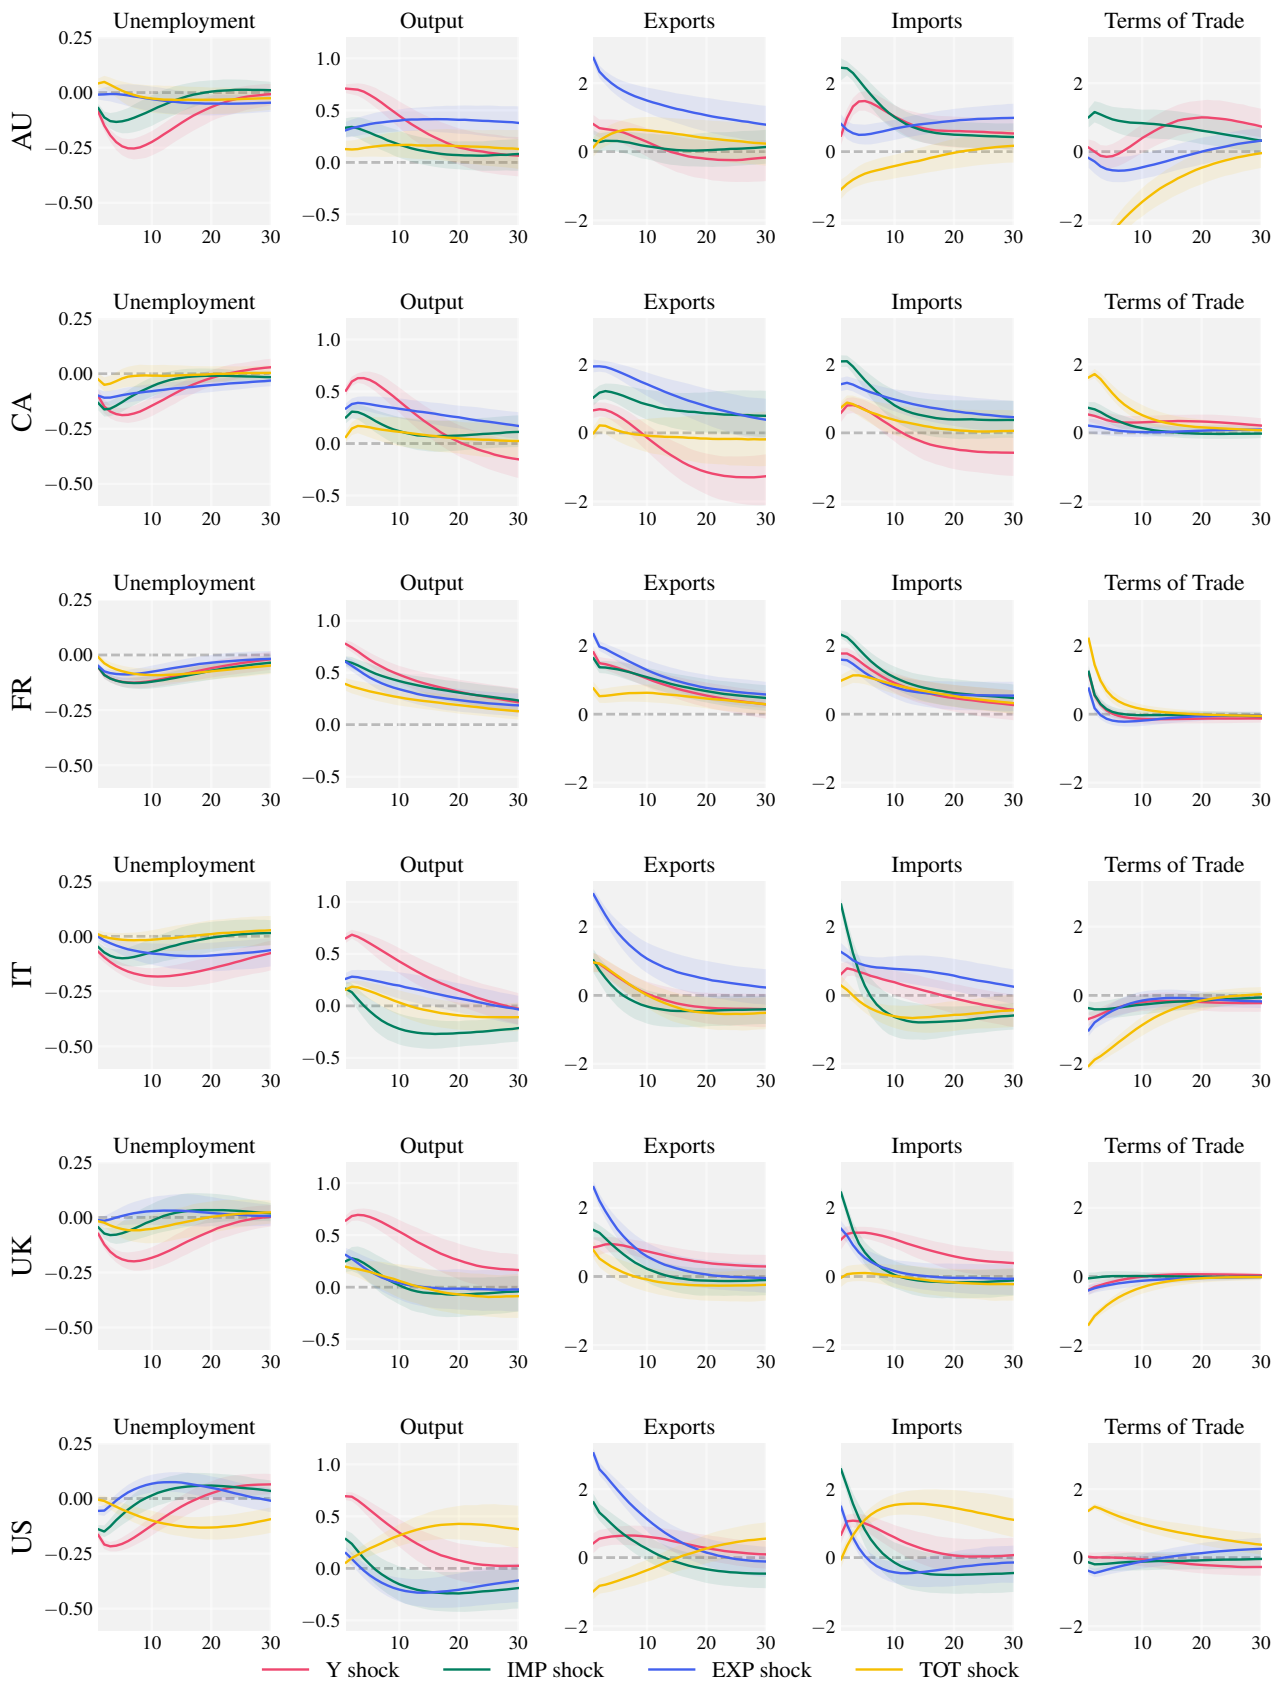
	u_t	Y_t	EXP_t	IMP_t	TOT_t
SOE					
EXP_t	0.17	0.49	1.00	0.49	0.24
IMP_t	0.48	0.49	0.49	1.00	0.16
TOT_t	0.11	0.30	0.24	0.16	1.00
US					
EXP_t	0.20	0.16	1.00	0.54	-0.32
IMP_t	0.52	0.39	0.54	1.00	-0.13
TOT_t	-0.05	0.01	-0.32	-0.13	1.00

Table 2.29: The various facets of the MBC shock, variance contributions (VAR)

	u	Y	EXP	IMP	TOT
AU					
Unemployment	73.7	48.5	5.1	21.2	5.2
Output	41.3	87.6	14.7	26.4	5.6
Exports	1.7	14.1	79.0	3.6	4.1
Imports	16.8	21.5	4.2	70.7	11.5
Terms of Trade	4.0	3.3	6.2	10.3	84.3
CA					
Unemployment	73.5	27.0	13.4	29.9	3.4
Output	29.8	74.6	24.7	19.0	6.1
Exports	8.4	20.8	73.5	30.8	1.9
Imports	23.8	18.9	30.5	70.7	16.6
Terms of Trade	3.8	6.0	3.6	14.1	92.5
FR					
Unemployment	81.7	27.4	18.4	23.5	4.6
Output	28.9	92.9	53.5	51.5	20.1
Exports	18.0	54.1	90.4	39.1	9.3
Imports	28.6	58.3	45.0	81.3	19.0
Terms of Trade	10.3	22.4	6.9	21.5	76.7
IT					
Unemployment	73.6	12.4	2.2	10.5	1.9
Output	17.7	74.0	13.9	8.7	9.7
Exports	3.2	12.7	81.0	12.2	16.9
Imports	13.9	12.3	13.3	65.3	5.1
Terms of Trade	1.9	9.0	15.7	5.0	86.8
UK					
Unemployment	67.1	35.7	4.0	15.5	5.1
Output	29.7	77.0	16.6	32.3	8.8
Exports	2.7	14.1	84.2	26.3	8.6
Imports	11.0	19.0	31.4	70.5	1.9
Terms of Trade	4.6	8.2	8.5	2.4	94.3
US					
Unemployment	76.2	44.7	5.3	19.6	1.5
Output	45.1	76.5	5.5	19.2	1.4
Exports	8.6	9.1	88.6	21.7	10.5
Imports	23.4	17.8	26.6	66.1	2.9
Terms of Trade	2.8	3.7	10.7	14.3	87.8

The rows correspond to different targets in the construction of the shock. The columns give the contributions of the constructed shock to the business-cycle volatility of the variables.

Figure 2.18: Comparison of MBC shock with import, export and terms of trade shocks, IRFs (VAR)



IRFs of selected variables to selected shocks. Horizontal axis: Time horizon in quarters. Vertical axis: Percent.

Chapter 3

Costs of Leaning against the Wind: Evidence from High-Frequency Identification

Abstract

How costly is it for central banks to lean against housing prices or credit aggregates using their interest rate instrument? We provide a measure of these costs by combining state-of-the-art estimation and identification methods with exogenous monetary policy surprises from Bauer and Swanson (2022) as an external instrument. We find that the Federal Reserve would need to accept a real GDP loss of 1% to bring down real housing prices by merely 4%. Moreover, a real GDP loss of 1% would allow reducing real credit by only about 1.5%. The ability of the central bank to correct disequilibria in asset price markets or credit markets thus appears limited and is associated with significant real economic costs. Our results are in line with the literature that advocates against leaning against the wind (LAW) policies.

Acknowledgment: Joint with Larissa Schwaller (University of Bern). We thank Luca Benati, Pierpaolo Benigno, Fabrice Collard, Harris Dellas, Daniel Kaufmann, Cyril Monnet, Dirk Niepelt, and Aubrey Poon for helpful comments and suggestions. Further, we are grateful for helpful comments by seminar and conference participants in the Macro Group Meeting and Macro PhD Seminar at the Department of Economics at the University of Bern.

3.1 Introduction

Since the Great Financial crisis, a discussion around leaning against the wind (LAW) policies has emerged. Suppose a central bank conducts a LAW policy. In that case, it pursues a somewhat tighter monetary policy, i.e., a higher policy rate, than what is consistent with flexible inflation targeting without taking financial stability effects into account (Svensson (2017)). On the one hand, LAW policies are costly because they lead to a weaker economy with a lower real Gross Domestic Product (GDP), higher unemployment, and a lower inflation rate. On the other hand, such policies can potentially reduce the probability and magnitude of crises.

The dominant view among economists is that the costs of LAW policies are likely to outweigh the benefits by a significant amount (see Svensson (2017), Schularick et al. (2021), Benati (2021), Benati (2022)). Additionally, the costs are expected to materialize with a very high probability, while the benefits are seen as highly uncertain. In contrast, there exists a minority position that central banks should pursue financial stability-oriented monetary policy, meaning that central banks should systematically lean against the wind (see BIS (2014), BIS (2016)). The advocates of such a monetary policy demand that central banks mainly pursue traditional objectives, i.e., concentrating on short-term output and inflation objectives, but deviate temporarily to avert a financial crisis when signs of financial imbalances emerge.

In our paper, we contribute to this discussion by providing a measure of the costs that central banks face when LAW using their interest rate instrument. We evaluate the costs in terms of relative real GDP losses. In particular, we estimate vector autoregression (VAR) models using frequentist and Bayesian methods and identify the monetary policy shocks following an external instrument strategy. We consider the method of Gertler and Karadi (2015) for the frequentist approach and the method of Arias, Rubio-Ramírez, et al. (2021) for the Bayesian approach. As external instruments, also called proxies, we use the orthogonalized monetary policy surprises (MPS) constructed by Bauer and Swanson (2022). MPS are high-frequency changes in interest rates around Federal Open Market Committee (FOMC) announcements. We look at two types of models, one focusing on housing prices and one focusing on credit variables. The former allows us to quantify the costs of leaning against housing prices, while the latter provides a measure of the costs of leaning against credit. We compute the costs by taking the ratio between the cumulative impulse response of either real housing prices or real credit and the cumulative impulse response of real GDP.¹ The ratio can be calculated for different horizons.

Our analysis aims to provide a new evaluation of the costs of LAW policies. We employ state-of-the-art methods to identify the monetary policy shock, which is a key ingredient for calculating these costs. In recent years, great advances have been made in monetary policy shock identification. Much of the literature has shifted away from using zero or sign restrictions

¹The considered impulse responses are the impulse responses to a monetary policy shock.

and towards strategies employing external instruments. The advantage of using the latter is that this method imposes less theoretically motivated restrictions. Specifically, the external instrument approach does not rely on taking a stance on the ordering of variables – as done by zero restrictions² – nor on imposing a particular theory on how a monetary policy shock should affect certain variables – as done with sign restrictions³. We employ the orthogonalized MPS of Bauer and Swanson (2022) as an instrument. These surprises differ from those previously utilized in the literature in that they explicitly correct certain endogeneity issues. Some recent studies have questioned the exogeneity of MPS because they found, for instance, that these surprises correlate with publicly available macroeconomic and financial data before the FOMC announcement. Bauer and Swanson (2022) account for this problem by orthogonalizing the surprises with respect to macroeconomic and financial data that pre-date the announcement.

We show that the impulse response functions (IRFs) to a monetary policy shock in both datasets largely follow conventional wisdom. Moreover, we demonstrate that the IRFs obtained with the frequentist approach are very similar to the ones obtained with the Bayesian approach. The main difference between the two are the confidence bands, which are substantially larger in the Bayesian setting. This observation leads us to evaluate the costs of LAW policies for the United States (US) based on the IRFs of the frequentist method.

Our evidence points towards an unfavorable outcome of LAW policies in terms of relative real GDP losses. The impact of an unexpected 25 basis point (bp) interest rate hike on real housing prices or real credit is not significantly larger than the impact on real GDP. In particular, the central bank would need to accept a real GDP loss of 1% to bring down real house prices by merely 4%. When it comes to real credit, the costs are even higher. For a real GDP loss of 1%, real credit could only be decreased by about 1.5%. Our results are in line with the dominant view in the literature that advocates against pursuing LAW policies.

The remainder of this paper is structured as follows. Section 3.2 discusses the relevant literature. Section 3.3 describes the data. Section 3.4 presents our estimation procedure, the identification strategy, and the instrument. In Section 3.5, we describe the results. Finally, Section 3.6 concludes.

3.2 Literature

This paper primarily contributes to the literature on LAW policies. The dominant view in this literature is that the costs of LAW policies are likely to exceed the benefits. Svensson (2017) proposes a simple framework for the cost-benefit analysis of LAW and finds that the costs of

²Zero restrictions assume, for example, that within the current month, monetary policy reacts to the observed output and inflation, but output and inflation do not react to monetary policy during the same time period.

³Sign restrictions impose, for instance, that a monetary policy shock affects output and inflation negatively two quarters after impact.

LAW policies exceed the benefits by a substantial margin. He argues that the empirical interest rate effects in terms of a lower probability and magnitude of crises are too small to match the costs. He also shows that the crisis loss and the cost of a crisis, i.e., the crisis loss minus the non-crisis loss, for a given magnitude of a crisis are higher with LAW than without. Svensson (2017) calls this last finding the second cost of LAW. Intuitively, if a crisis occurs when the central bank is pursuing a LAW policy (and the economy is therefore weaker), then – given a certain magnitude of the crisis – the economy will also be more vulnerable in the crisis.

Schularick et al. (2021) conduct a state-dependent analysis and focus only on periods of financial booms. Based on observations of boom periods for 17 developed countries, they identify the causal effects of a monetary policy shock. They show that discretionary LAW policies during financial booms are more likely to trigger rather than prevent crises, i.e., LAW policies increase crisis probability. Additionally, discretionary LAW policies do not systematically reduce crisis severity. They employ local projections combined with an instrumental variable to identify the monetary policy shocks.

Benati (2021) explores the trade-off faced by central banks between leaning against housing prices and stabilizing real economic activity. He uses monetary VARs and proposes two methods for evaluating the trade-off. The first method consists of identifying monetary policy shocks along the lines of Arias, Caldara, et al. (2019) by imposing zero or sign restrictions on the contemporaneous response of the policy rate to the other series in the VAR. Moreover, he adds additional constraints on the signs of the reactions of the policy rate, commodity prices, and monetary aggregates to a monetary shock, following the work of Uhlig (2005). He finds that in the US, Canada, and the United Kingdom (UK), the effect of a monetary policy shock on real housing prices is about three to five times as large as that on real GDP. At first glance, these trade-offs seem relatively favorable. However, when considering the large difference in magnitude between housing price growth and GDP growth, significant GDP losses would need to be accepted to bring down housing prices by a sufficient amount. For the second method and based on the identified VAR, Benati (2021) computes the trade-off associated with weakly but systematically LAW. Following the strategy of Leeper and Zha (2003), he computes the counterfactual of a marginally more aggressive policy rule towards fluctuations in specific variables – here, the ratio between housing prices and rents. The results indicate that for the three considered countries, the impact on real housing prices is about three times as large as the one on real GDP.

Benati (2022) conducts a similar analysis as Benati (2021) but focuses on the trade-off between leaning against credit and stabilizing real economic activity. He runs his analysis for ten countries. The evidence points towards an unfavorable trade-off for all countries. The effect of a monetary policy shock on either real credit or credit leverage is not substantially higher than its effect on real GDP. Benati (2022) also finds that modest systematic interventions by

the Federal Reserve during the years before the financial crisis would have led to a two and a half percent decrease in credit leverage per one percent decline in real GDP.

The minority view in this literature is that central banks should systematically LAW. The representatives of this view are mainly economists at the Bank of International Settlements (BIS). In the annual reports of 2014 and 2016 (BIS (2014), BIS (2016)), they argue in favor of a monetary policy that is ready to tighten whenever financial imbalances show signs of building up, even if inflation seems to be under control in the short term. The idea is to pursue a more systematic and structured way of assessing the risks that slower-moving financial cycles pose to macroeconomic stability, inflation, and the effectiveness of policy tools.

Our work also relates to the line of research on monetary policy shock identification. In her work, Ramey (2016) highlights that the estimated dynamic responses to monetary policy shocks can be susceptible to the choice of the empirical specification, sample, and instrument. Moreover, she explains the issue of possible foresight about monetary policy changes. The problem of foresight relates to the growing importance of forward guidance in monetary policy and consists of the fact that many changes in policy rates may be anticipated. In recent years, the approach of identifying monetary policy shocks using Proxy-Structural-VARs (Proxy-SVARs) received considerable attention. This framework has been developed by Stock and Watson (2012) and Mertens and Ravn (2013). Gertler and Karadi (2015) then applied this approach to monetary policy shocks and used high-frequency instruments (HFI) as proxies. They consider MPS as the HFI. MPS measure high-frequency changes in different types of interest rates around FOMC announcements. The critical advantage of HFI is that they provide a way of dealing with foresight issues. Nevertheless, while the general idea behind the methodology of Gertler and Karadi (2015) was compelling, their instruments suffered severe problems. Bauer and Swanson (2022), along with other studies, question the exogeneity of the MPS. For instance, the surprises used by Gertler and Karadi (2015) seem to be correlated with economic and financial information that pre-dates the FOMC announcements. Hence, Bauer and Swanson (2022) correct for this problem by orthogonalizing the surprises with respect to macroeconomic and financial data that pre-date the announcement. Further, Arias, Rubio-Ramírez, et al. (2021) propose a way of implementing the Proxy-SVAR method in a Bayesian framework. In their paper, they do not specifically focus on the identification of monetary policy shocks. However, by using the MPS of Bauer and Swanson (2022) as an instrument, their method can be easily used for this application.

Our contribution consists in combining the literature on LAW policies with the current literature on the identification of monetary policy shocks. We employ both the methods of Gertler and Karadi (2015) and Arias, Rubio-Ramírez, et al. (2021) and connect them with the MPS of Bauer and Swanson (2022). Identifying monetary policy shocks based on state-of-the-art methods provides a highly improved measure of the dynamic responses to these shocks. As

a result, the estimated costs of pursuing LAW policies become more accurate.

3.3 Data

We use monthly US data from 1973:1 to 2019:12. The start of the sample is determined by the earliest availability of the excess bond premium (EBP). We end our sample in 2019:12 to exclude the COVID-19 pandemic and its aftermath. We consider two datasets, one focusing on housing prices and one on credit aggregates. The variables common to both datasets are the logarithms of real GDP⁴, the core Personal Consumption Expenditures (PCE) deflator⁵, an index of commodity prices⁶ the real effective exchange rate (REER); the EBP by Gilchrist and Zakrajšek (2012)⁷, and the two-year Treasury yield. These six variables constitute our baseline dataset. We include the EBP because Caldara and Herbst (2019) find it to be necessary to identify monetary policy shocks correctly. To be precise, they show that central banks react systematically to changes in corporate credit spreads. The failure to account for this systematic reaction leads to an attenuation in the response of all variables to a monetary policy shock. Following the discussion in Gertler and Karadi (2015), we consider the two-year Treasury yield instead of the Federal Funds rate. Unlike the latter, the former was essentially unconstrained during the zero lower bound period in the US between 2009 and 2015 and is thus seen as a better measure of the stance of monetary policy. Additionally, the advantage of using a government bond rate as the policy indicator is that its innovations do not only capture traditional monetary policy shocks, i.e., surprises in the current federal funds rate but also shocks to forward guidance. Moreover, we include the REER because the relationship between exchange rates and monetary policy is a classic topic in monetary economics.⁸ In addition to the common variables, the housing dataset includes the logarithms of a house price index⁹, a rent price index¹⁰, employees working in construction, real estate loans; and the 30-

⁴We use real GDP rather than industrial production (IP) as the latter has decreased in importance in the US over time. Specifically, US IP is only about a fifth of GDP, so what happens to IP is not necessarily fully representative of what happens in the economy at large.

⁵The core PCE deflator excludes items that tend to be very volatile, like food and energy prices. Even if these goods often make up an essential part of a consumer's budget, focusing on core inflation measures proves helpful for a central bank in assessing inflation trends.

⁶We use the log of the Commodity Research Bureau's monthly index of commodity prices. We include commodity prices because the response of this variable is often of interest and has been considered by many previous papers that identify monetary policy shocks.

⁷Gilchrist and Zakrajšek (2012) construct a corporate bond credit spread index – the so-called GZ credit spread, which is based on a large micro-level dataset. They then decompose the GZ credit spread into two parts: one part capturing the systematic movements in default risk of individual firms and a residual component – the EBP. The EBP can be interpreted as the variation in the pricing of default risk, meaning it is a measure of the tightness of financial conditions.

⁸See e.g. the classic Mundell-Fleming framework and Eichenbaum and Evans (1995).

⁹We consider the median sales prices of new homes sold in the US.

¹⁰For the rent series, we use a shelter component of the consumer price index.

year fixed rate mortgage average. We include employees working in construction because it is an important cyclical variable that captures the state of economic activity in the construction sector. The credit dataset comprises the logarithms of real estate loans and non real estate loans¹¹; M1 velocity¹², and the ratio of the monetary aggregates M1 over M2. We include M1 velocity and the ratio of M1 over M2 because broad money and credit are essentially two faces of the same coin. The main activity of commercial banks consists of taking on deposits on the liability side and extending credit on the asset side. Hence, to fully characterize monetary policy transmission to credit aggregates, we also have to include monetary aggregates in our dataset. Finally, we demean all series. The data are described in more detail in Appendix A.

3.4 Empirical Strategy

In this paper, we use Proxy-SVAR methods to estimate and identify monetary policy shocks. We consider two types of Proxy-SVAR approaches: the one by Gertler and Karadi (2015), which is based on frequentist estimation methods, and the one by Arias, Rubio-Ramírez, et al. (2021), which uses Bayesian methods. These types of SVARs require a proxy to identify the shock of interest. We employ the MPS published by Bauer and Swanson (2022) as the external instrument. These surprises are orthogonalized with respect to macroeconomic and financial data that pre-dates the FOMC announcements, which significantly improves the exogeneity dimension of this instrument. This section briefly discusses the methods of Gertler and Karadi (2015) and Arias, Rubio-Ramírez, et al. (2021) as well as the MPS of Bauer and Swanson (2022) to make the present paper self-contained.

3.4.1 Estimation and Identification

Frequentist Approach

For the frequentist approach, we follow the method by Gertler and Karadi (2015). To describe their estimation and identification strategy, we start by considering the following structural VAR:

$$\mathbf{y}'_t \mathbf{A}_0 = \sum_{\ell=1}^p \mathbf{y}'_{t-\ell} \mathbf{A}_\ell + \mathbf{c} + \boldsymbol{\varepsilon}'_t, \quad (3.4.1)$$

where \mathbf{y}_t is an $n \times 1$ vector of observables, $\boldsymbol{\varepsilon}_t$ is an $n \times 1$ vector of white noise structural shocks, and \mathbf{c} is an $1 \times n$ vector of constants. \mathbf{A}_i is an $n \times n$ matrix for $0 \leq i \leq p$ and \mathbf{A}_0 is invertible.

¹¹For non real estate loans, we use the sum of consumer loans, commercial and industrial loans, and all other loans and leases from all commercial banks.

¹²M1 velocity is computed as the ratio of nominal GDP over the monetary aggregate M1. Following Benati et al. (2021), we calculate M1 as the sum of M1 and Money Market Deposit Accounts (MMDA) as this restores the long-run equilibrium relationship between M1 velocity and interest rates.

When multiplying both sides with \mathbf{A}_0^{-1} , the reduced-form VAR representation follows:

$$\mathbf{y}'_t = \sum_{\ell=1}^p \mathbf{y}'_{t-\ell} \mathbf{B}_\ell + \mathbf{c} \mathbf{A}_0^{-1} + \mathbf{u}'_t, \quad (3.4.2)$$

with \mathbf{u}_t being the reduced-form VAR residuals, $\mathbf{B}_\ell = \mathbf{A}_\ell \mathbf{A}_0^{-1}$, and $\mathbb{E}[\mathbf{u}_t \mathbf{u}'_t] = \mathbf{\Sigma}$ for some positive definite matrix $\mathbf{\Sigma}$. The VAR residuals are linear combinations of the underlying structural shocks, namely

$$\mathbf{u}'_t = \boldsymbol{\varepsilon}'_t \mathbf{S}'. \quad (3.4.3)$$

It follows that $\mathbf{S}' = \mathbf{A}_0^{-1}$ and $\mathbb{E}[\mathbf{u}_t \mathbf{u}'_t] = \mathbb{E}[\mathbf{S} \mathbf{S}'] = \mathbf{\Sigma}$.

Let us then define $y_t^p \in \mathbf{y}_t$ to be the monetary policy indicator, i.e., the variable for which the exogenous variation is due to the monetary policy shock ε_t^p . To estimate the impulse responses to a monetary policy shock, we need to estimate the equation

$$\mathbf{y}'_t = \sum_{\ell=1}^p \mathbf{y}'_{t-\ell} \mathbf{B}_\ell + \mathbf{c} \mathbf{A}_0^{-1} + \varepsilon_t^p \mathbf{s}', \quad (3.4.4)$$

where \mathbf{s} is the column of \mathbf{S} associated with the effects of ε_t^p . Because we are only interested in the impulse responses to a monetary policy shock, it is sufficient to identify \mathbf{s} and not the entire matrix \mathbf{S} . To achieve the identification of \mathbf{s} , we use an external instruments strategy.

We define \mathbf{m}_t to be the $k \times 1$ vector of instruments. $\boldsymbol{\varepsilon}_t^q$ is a vector of structural shocks other than the monetary policy shock. For \mathbf{m}_t to be a valid set of instruments, the exogeneity and relevance conditions must be satisfied:

$$\begin{aligned} \mathbb{E}[\mathbf{m}_t \varepsilon_t^p] &\neq 0 \\ \mathbb{E}[\mathbf{m}_t (\boldsymbol{\varepsilon}_t^q)'] &= \mathbf{0}, \end{aligned} \quad (3.4.5)$$

meaning that the instruments have to be correlated with the monetary policy shock ε_t^p but orthogonal to any other structural shock in $\boldsymbol{\varepsilon}_t^q$. We will always focus on a single instrument in the following application, namely the MPS of Bauer and Swanson (2022).

The identification of \mathbf{s} works as follows: First, we estimate the VAR using least squares estimation and get the reduced-form residuals \mathbf{u}_t . These residuals can then be split up into u_t^p , the residual associated with the equation of the policy indicator, and \mathbf{u}_t^q , the residuals of all other equations $q \neq p$. Moreover, we determine $s^p \in \mathbf{s}$ to be the response of u_t^p to a unit increase in ε_t^p . Similarly, $\mathbf{s}^q \in \mathbf{s}$ is the response of \mathbf{u}_t^q to an increase of ε_t^p by one unit. Second, we perform a two-stage least squares regression. In the first stage, we regress u_t^p on the instrument \mathbf{m}_t . Consequently, the variation in the fitted value \hat{u}_t^p is only due to the monetary policy shock

ε_t^p . In the second stage, we regress \mathbf{u}_t^q on \hat{u}_t^p :

$$\mathbf{u}_t^q = \frac{\mathbf{s}^q}{s^p} \hat{u}_t^p + \boldsymbol{\xi}_t. \quad (3.4.6)$$

This regression yields a consistent estimate of $\frac{\mathbf{s}^q}{s^p}$ because \hat{u}_t^p is uncorrelated with the error term $\boldsymbol{\xi}_t$. An estimate for s^p can be obtained from the estimated variance-covariance matrix $\boldsymbol{\Sigma}$. In the next step, \mathbf{s}^q can be computed. Based on the estimates of s^p , \mathbf{s}^q , and the VAR coefficients ($\mathbf{B}_{\ell s}$), we can calculate the impulse responses of all variables in \mathbf{y}_t to a monetary policy shock ε_t^p . To get the confidence bands around the point estimates, we use bootstrapping methods, with 10,000 bootstrap replications.¹³ We set the lag order to $p = 12$.

Bayesian Approach

For the Bayesian approach, we use the method of Arias, Rubio-Ramírez, et al. (2021). We write the structural VAR as

$$\tilde{\mathbf{y}}_t' \tilde{\mathbf{A}}_0 = \sum_{\ell=1}^p \tilde{\mathbf{y}}_{t-\ell}' \tilde{\mathbf{A}}_{\ell} + \tilde{\mathbf{c}} + \tilde{\boldsymbol{\varepsilon}}_t', \quad (3.4.7)$$

where $\tilde{\mathbf{y}}_t' = [\mathbf{y}_t' \mathbf{m}_t']$ with dimensions $1 \times \tilde{n}$ and $\tilde{n} = n + k$. $\tilde{\mathbf{A}}_i$ is an $\tilde{n} \times \tilde{n}$ matrix for $0 \leq i \leq p$ with $\tilde{\mathbf{A}}_0$ being invertible, $\tilde{\mathbf{c}}$ is a $1 \times \tilde{n}$ row vector, and $\tilde{\boldsymbol{\varepsilon}}_t$ is conditionally standard normal.

By defining $\tilde{\mathbf{x}}_t' = [\tilde{\mathbf{y}}_{t-1}' \cdots \tilde{\mathbf{y}}_{t-p}' 1]$ and $\tilde{\mathbf{A}}_+' = [\tilde{\mathbf{A}}_1' \cdots \tilde{\mathbf{A}}_p' 1]$, we can rewrite (3.4.7) as

$$\tilde{\mathbf{y}}_t' \tilde{\mathbf{A}}_0 = \tilde{\mathbf{x}}_t' \tilde{\mathbf{A}}_+ + \tilde{\boldsymbol{\varepsilon}}_t'. \quad (3.4.8)$$

Let $\tilde{\boldsymbol{\varepsilon}}_t' = [\boldsymbol{\varepsilon}_t' \boldsymbol{\nu}_t']$, where the $n \times 1$ vector $\boldsymbol{\varepsilon}_t$ contains the structural shocks and the $k \times 1$ vector $\boldsymbol{\nu}_t$ the other shocks that affect the proxies.¹⁴ In a Proxy-SVAR, \mathbf{y}_t evolves according to

$$\mathbf{y}_t' \mathbf{A}_0 = \mathbf{x}_t' \mathbf{A}_+ + \boldsymbol{\varepsilon}_t', \quad (3.4.9)$$

with $\mathbf{x}_t' = [\mathbf{y}_{t-1}' \cdots \mathbf{y}_{t-p}' 1]$ and $\mathbf{A}_+' = [\mathbf{A}_1' \cdots \mathbf{A}_p' \mathbf{c}']$. \mathbf{A}_i is an $n \times n$ matrix for $0 \leq i \leq p$, \mathbf{A}_0 is invertible, and \mathbf{c} is a $1 \times n$ row vector. $(\mathbf{A}_0, \mathbf{A}_+)$ are called the structural SVAR parameters. Given that \mathbf{y}_t has to follow (3.4.9), we need to impose the following zero restrictions on $\tilde{\mathbf{A}}_i$:

$$\tilde{\mathbf{A}}_i = \begin{pmatrix} \mathbf{A}_i & \boldsymbol{\Gamma}_{i,1} \\ \mathbf{0}_{k \times n} & \boldsymbol{\Gamma}_{i,2} \end{pmatrix}, \quad (3.4.10)$$

where $\boldsymbol{\Gamma}_{i,1}$ is $n \times k$ and $\boldsymbol{\Gamma}_{i,2}$ is $k \times k$ for $0 \leq i \leq p$. Arias, Rubio-Ramírez, et al. (2021) refer to these zero restrictions on $\tilde{\mathbf{A}}_i$ as block restrictions. Moreover, they call (3.4.8) together with the

¹³We are using the wild bootstrap procedure of Mertens and Ravn (2013) and Gertler and Karadi (2015).

¹⁴ $\tilde{\boldsymbol{\varepsilon}}_t$ is conditionally standard normal. Hence, $\boldsymbol{\varepsilon}_t$ and $\boldsymbol{\nu}_t$ are uncorrelated.

block restrictions the structural parametrization of the Proxy-SVAR and $(\tilde{\mathbf{A}}_0, \tilde{\mathbf{A}}_+)$, conditional on the block restrictions being satisfied, the Proxy-SVAR structural parameters.

In our application, we use the MPS of Bauer and Swanson (2022) as a single proxy. Thus, we have $k = 1$. To identify the monetary policy shock, we have to assume that the proxy is correlated with the monetary policy shock $\varepsilon_t^p \in \boldsymbol{\varepsilon}_t$ (relevance condition) and uncorrelated with the remaining $n - 1$ structural shocks (exogeneity restrictions). Let us multiply (3.4.8) by $\tilde{\mathbf{A}}_0^{-1}$ and by $\mathbf{J}' = [\mathbf{0}_{1 \times n} \ \mathbf{I}_1]'$. We then obtain $\mathbf{m}'_t = \tilde{\mathbf{y}}'_t \mathbf{J}' = \tilde{\mathbf{x}}'_t \tilde{\mathbf{A}}_+ \tilde{\mathbf{A}}_0^{-1} \mathbf{J}' + \tilde{\boldsymbol{\varepsilon}}'_t \tilde{\mathbf{A}}_0^{-1} \mathbf{J}'$. It follows that

$$\mathbf{E}[\mathbf{m}_t \boldsymbol{\varepsilon}'_t] = \mathbf{E}[\mathbf{m}_t \tilde{\boldsymbol{\varepsilon}}'_t \mathbf{L}'] = \mathbf{J}(\tilde{\mathbf{A}}_0^{-1})' \mathbf{L}', \quad (3.4.11)$$

where $\mathbf{L} = [\mathbf{I}_n \ \mathbf{0}_{n \times 1}]$. Hence, the exogeneity restriction requires that the first $n - 1$ columns of matrix $\mathbf{J}(\tilde{\mathbf{A}}_0^{-1})' \mathbf{L}'$ must be zero. Additionally, the relevance condition is satisfied if the last column of $\mathbf{J}(\tilde{\mathbf{A}}_0^{-1})' \mathbf{L}'$ is non-singular.

In the first step of the algorithm, we want to create a proposal for the desired distribution over the structural parameters conditional on the exogeneity restrictions, the γ -relevance condition, and any additional zero and sign restrictions. The traditional approach of generating such a proposal would be to map independent draws from the orthogonal reduced-form parametrization conditional on the restrictions into the structural parametrization of the SVAR conditional on the restrictions. In their paper, Arias, Rubio-Ramírez, et al. (2021) show that the traditional approach cannot be applied in this context because alone the number of zero restrictions implied by the block restrictions is too large. Hence, we use the orthogonal triangular-block parametrization of the Proxy-SVAR:

$$\tilde{\mathbf{y}}'_t \tilde{\mathbf{A}}_0 = \tilde{\mathbf{x}}'_t \tilde{\mathbf{A}}_+ + \tilde{\mathbf{u}}'_t, \quad (3.4.12)$$

where $\tilde{\mathbf{u}}'_t = \tilde{\boldsymbol{\varepsilon}}'_t \mathbf{Q}'$, with $\mathbf{Q} = \text{diag}(\mathbf{Q}_1, \mathbf{Q}_2)$.¹⁵ The zero restrictions pertaining to the exogeneity condition are linear restrictions on the columns of the orthogonal matrix \mathbf{Q}_1 . We make independent draws from the orthogonal triangular-block parameters $(\tilde{\mathbf{A}}_0, \tilde{\mathbf{A}}_+, \mathbf{Q}_1, \mathbf{Q}_2)$ and map them into Proxy-SVAR structural parameters by

$$(\tilde{\mathbf{A}}_0, \tilde{\mathbf{A}}_+, \mathbf{Q}_1, \mathbf{Q}_2) \xrightarrow{f} \underbrace{(\tilde{\mathbf{A}}_0 \text{diag}(\mathbf{Q}_1, \mathbf{Q}_2))}_{\tilde{\mathbf{A}}_0}, \underbrace{(\tilde{\mathbf{A}}_+ \text{diag}(\mathbf{Q}_1, \mathbf{Q}_2))}_{\tilde{\mathbf{A}}_+}. \quad (3.4.13)$$

To impose the relevance condition, only those draws of $(\tilde{\mathbf{A}}_0, \tilde{\mathbf{A}}_+)$ are retained, which satisfy the γ -relevance condition. In Arias, Rubio-Ramírez, et al. (2021), they choose a γ -relevance criterion of 0.2. For our application, we use a criterion of $\gamma = 0.1$. This implies that at least 10% of the variance of the instrument must be related to the underlying shock of interest.

¹⁵Like $\tilde{\boldsymbol{\varepsilon}}_t$, the innovations $\tilde{\mathbf{u}}_t$ are conditionally standard normal.

Draws that do not satisfy this criterion are rejected.

The independent draws of $(\tilde{\mathbf{A}}_0, \tilde{\mathbf{A}}_+)$ are generated using the Gibbs sampler from Waggoner and Zha (2003). This Gibbs sampler allows us to draw from a normal-generalized normal (NGN) distribution characterized by $(\tilde{\nu}, \tilde{\Phi}, \tilde{\Psi}, \tilde{\Omega})$, subject to linear restrictions.¹⁶ These draws are then mapped into Proxy-SVAR structural parameters to get a proposal for the desired distribution. The obtained distribution of $(\tilde{\mathbf{A}}_0, \tilde{\mathbf{A}}_+)$ is proportional to an $\text{NGN}(\tilde{\nu}, \tilde{\Phi}, \tilde{\Psi}, \tilde{\Omega})$. We weigh the proposal draws using an importance sampler to get to our target distribution. The importance weights are computed as

$$w_i = \frac{\text{NGN}_{(\tilde{\nu}, \tilde{\Phi}, \tilde{\Psi}, \tilde{\Omega})}(\tilde{\mathbf{A}}_0, \tilde{\mathbf{A}}_+)}{p(\tilde{\mathbf{A}}_0, \tilde{\mathbf{A}}_+)}. \quad (3.4.14)$$

If not specified otherwise, we generate 40,000 independent draws and set the lag order to $p = 12$. Compared to the frequentist method, a clear drawback of the Bayesian approach is its computational costliness.¹⁷

3.4.2 Instrument

The Federal funds futures (FFF) market allows market participants to hedge against market fluctuations in the Federal Funds Rate (FFR). On any current day, the FFF market continuously reflects the market's expectations of the average FFR over the remainder of the month. Suppose, due to an FOMC announcement, there is an upward or downward revision in the current month's future rate. In that case, this is evidence that market participants were surprised by the policy announcement. The construction of MPS builds on this idea and measures intraday price changes in FFF contracts in a narrow window (typically 30 minutes) around FOMC announcements. The use of a narrow time window around the announcement aims to rule out reverse causality and to ensure that the change in the FFF market is due to the announcement and not due to other economic events.

To construct their MPS, Bauer and Swanson (2022) follow Gürkaynak et al. (2005) and use the change in a range of futures contracts, namely the change in the first four quarterly Eurodollar futures contracts. They then extract the first principal component of the changes in these four futures contracts around FOMC announcements. This approach allows to capture the surprise about the current and expected future path of the FFR. Using this method, they get the so-called 'raw' MPS. They further illustrate that these raw MPS are systematically

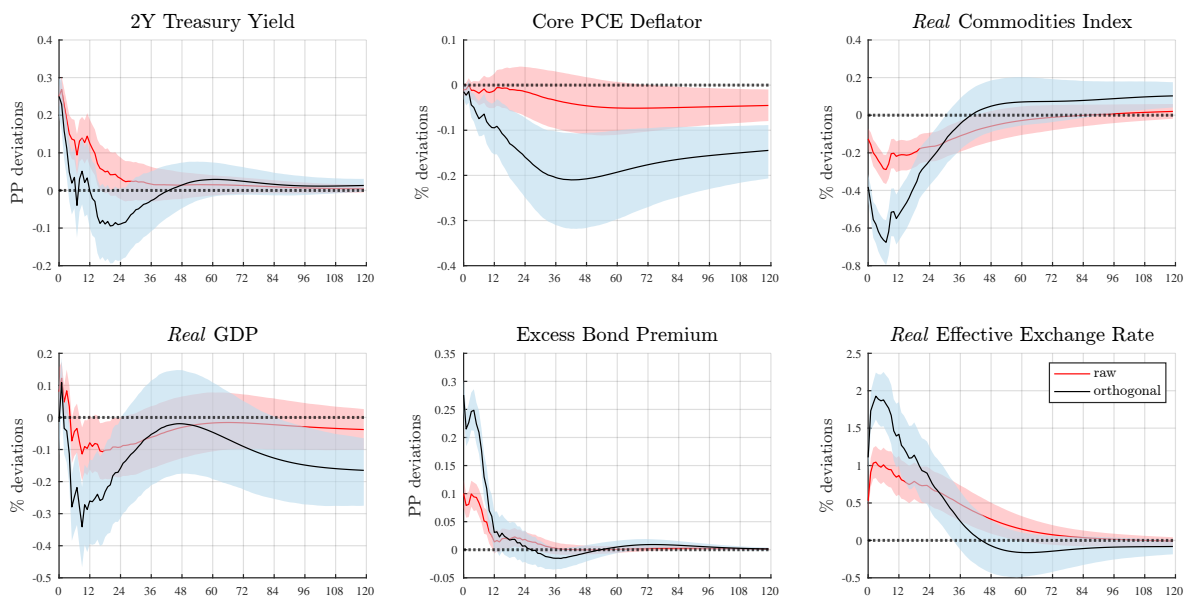
¹⁶Generally, $(\tilde{\nu}, \tilde{\Phi}, \tilde{\Psi}, \tilde{\Omega})$ can be set equal to $(\tilde{\nu}, \tilde{\Phi}, \tilde{\Psi}, \tilde{\Omega})$. However, sometimes this choice can lead to small effective sample sizes in the importance sampler. Arias, Rubio-Ramírez, et al. (2021) suggest a specific choice of $(\tilde{\nu}, \tilde{\Phi}, \tilde{\Psi}, \tilde{\Omega})$ that can avoid this loss of efficiency.

¹⁷The main expense of the Arias, Rubio-Ramírez, et al. (2021) algorithm is computing the importance weights.

correlated with macroeconomic and financial data available prior to an announcement. This predictability violates the exogeneity assumption in both the frequentist and Bayesian approach (see Equations (3.4.5) and (3.4.11) respectively) and calls into question the suitability of MPS in high-frequency instrumental variable regressions and the thereby obtained results.

Bauer and Swanson (2022) address this exogeneity issue by removing the component of the raw MPS that is correlated with economic and financial data available before the FOMC announcement. Specifically, their orthogonalized MPS are the residuals from regressing the raw MPS on nonfarm payrolls surprises, employment growth, the S&P 500, commodity prices, the slope of the yield curve, and the Treasury skewness. To illustrate how the explicit treatment of MPS predictability affects the results, we compare the IRFs to a 25bp monetary policy shock using raw and orthogonalized MPS. We compute the IRFs for our baseline model consisting of the two-year Treasury yield, the core PCE deflator, commodity prices, real GDP, the EBP, and the REER. Figure 3.1 provides the comparison for the frequentist approach. The same comparison for the Bayesian approach can be found in Appendix 3.B.1, Figure 3.4.

Figure 3.1: IRFs with raw versus orthogonalized surprises



Note: The solid black lines and the blue shaded areas are the point estimates and the 16-84 percentiles using the orthogonal MPS. The solid red lines and the red shaded areas are the point estimates and 16-84 percentiles obtained with the raw MPS. IRFs are normalized to a 25bp increase in the two-year Treasury yield. Horizontal axis: time horizon in months. Results are based on the frequentist approach.

In line with the results of Bauer and Swanson (2022), we find for both methods that using the orthogonalized MPS leads to a larger reaction of the price level, real commodity prices, real output, and the EBP. Additionally, the real appreciation of the USD is stronger when using the orthogonalized instead of the raw surprises. Moreover, in the frequentist approach, the price level and real output responses are only significantly negative when using the adjusted

MPS. The empirical pattern obtained with orthogonalized MPS greatly improves the puzzling responses of the raw MPS, where neither the core PCE deflator nor real GDP reacts significantly negatively to a contractionary monetary policy shock. These results motivate us to follow the ‘best practice’ recommendation of Bauer and Swanson (2022) and only use orthogonalized MPS in the remainder of the paper. The orthogonalized MPS of Bauer and Swanson (2022) are available for the period from 1988:1 to 2019:12.¹⁸

3.5 Evidence

The main objective of this paper is to provide a measure of the costs a central bank faces when it aims to lean against housing prices or credit aggregates using its interest rate instrument. The computation of these costs requires us first to obtain the IRFs of real housing prices, real credit, and real GDP to a monetary policy shock. We describe the IRFs corresponding to our housing and credit datasets in the following section. The impulse responses are always normalized to a 25bp increase in the two-year Treasury yield. In a second step, we present our measure of the costs associated with LAW policies, based on the ratio between the cumulative IRF of either real housing prices or real credit and the cumulative IRF of real GDP. We calculate these ratios three, four, and five years after the interest rate hike.

3.5.1 Housing Prices

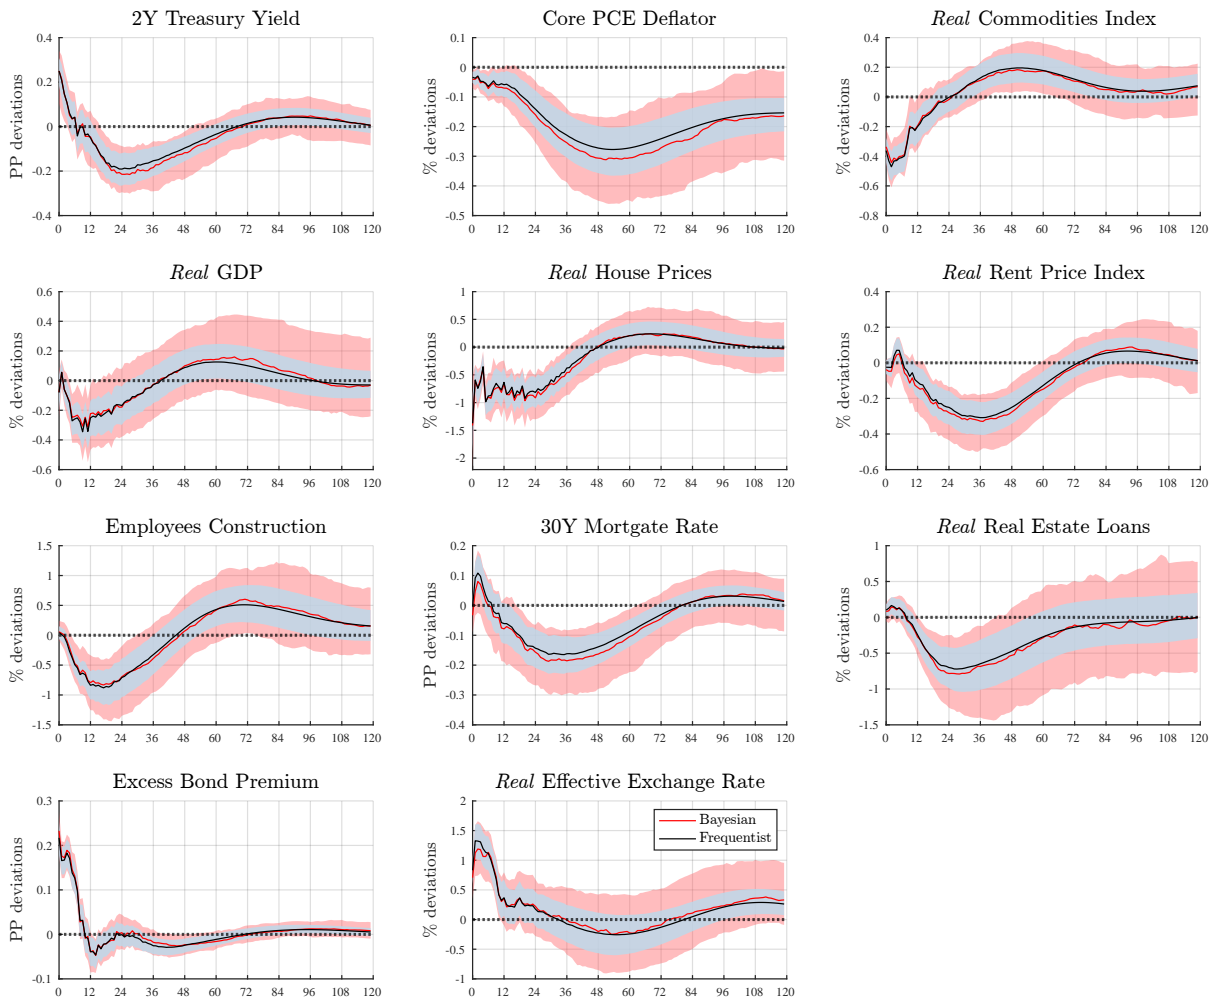
Figure 3.2 illustrates the effect of a 25bp increase in the two-year Treasury yield on macroeconomic variables, including house prices, rent prices, employees in the construction sector, the 30-year mortgage rate and real estate loans. The solid black line and the blue shaded area are the point estimate and the confidence bands resulting from the frequentist approach. The solid red line and the red shaded area show the median and the confidence bands obtained with the Bayesian approach. The two sets of IRFs are highly similar. The main difference between the two is the confidence bands, which are substantially larger in the Bayesian setting. This observation leads us to focus on the IRF from the frequentist approach.

The reactions of the price level and real GDP are in line with standard macroeconomic models. The price level decreases significantly already on impact, while real GDP does so within about three months after the interest rate hike. The reactions of the EBP and the REER are also standard. A contractionary monetary policy shock leads to a tightening of financial conditions and a real appreciation of the USD. Concerning real commodity prices, we

¹⁸For both the frequentist and Bayesian method, we estimate the VAR using the long sample from 1973:1 to 2019:12. For the identification, we then use the shorter sample based on the availability of the instrument. In the algorithm of Arias, Rubio-Ramírez, et al. (2021), we implement the two different sample lengths for the estimation and identification by setting the observations for the MPS from 1973:1 to 1987:12 equal to zero.

find that they fall on impact and stay negative for about 20 months. After that, real commodity prices turn positive for several years before reverting back to zero. The fast rebound and the medium-run increase of real commodity prices differ from conventional wisdom. We take this as a motivation in subsequent research to impose additional sign restrictions and to prevent this puzzling response. Except for real commodity prices, we find that adding more variables does not change the dynamics of the baseline variables, as observed in Figure 3.1.

Figure 3.2: IRFs to a monetary policy shock in the housing dataset



Note: The solid black lines and the blue shaded areas are the point estimates and the 16-84 percentiles resulting from the frequentist approach. The solid red lines and the red shaded areas are the medians and the 16-84 percentiles obtained with the Bayesian approach. IRFs are normalized to a 25bp increase in the two-year Treasury yield. Horizontal axis: time horizon in months.

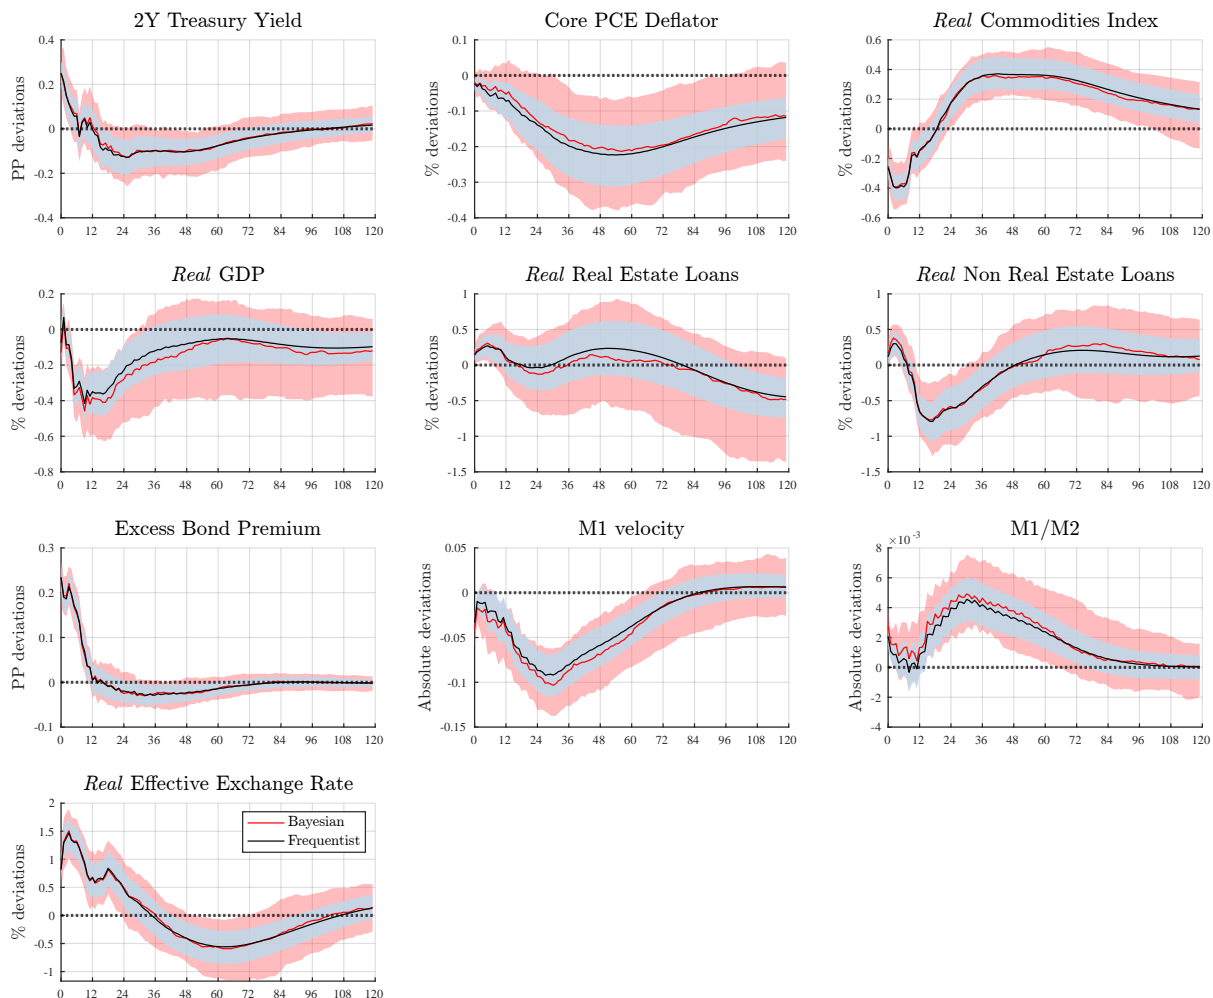
Concerning housing market variables, we find that real house prices fall significantly on impact and take about three to four years to return to their initial level. Real rent prices display more inertia and decrease only after about one year. They remain persistently below their initial values for about five years. Employees in construction do not react on impact,

but the response becomes negative already after roughly two months. The recovery back to the initial level takes about three years. Lastly, for the long-term mortgage rate, we see an initial increase which – after about one year and a half – is followed by a significant decline. Altogether, we find that the reaction to the contractionary monetary policy shock is in line with conventional wisdom because real output and prices fall, and the housing market is negatively affected.

3.5.2 Credit

Figure 3.3 displays the effect of a 25bp increase in the two-year Treasury yield on macroeconomic variables, including real estate and non real estate credit aggregates, M1 velocity, and the M1 over M2 ratio.

Figure 3.3: IRFs to a monetary policy shock in the credit dataset



Note: The solid black lines and the blue shaded areas are the point estimates and the 16-84 percentiles resulting from the frequentist approach. The solid red lines and the red shaded areas are the medians and the 16-84 percentiles obtained with the Bayesian approach. IRFs are normalized to a 25bp increase in the two-year Treasury yield. Horizontal axis: time horizon in months.

We focus again on the IRFs obtained with the frequentist method. We find the same picture as in the housing dataset regarding the price level, real commodity prices, real GDP, the EBP, and the REER. For credit aggregates, we observe a conflicting reaction between non real estate loans and real estate loans. Non real estate loans fall after one year and then rebound, whereas real estate loans hardly react significantly negatively or positively. In contrast, real estate loans decrease significantly in the housing dataset after about one year. To investigate this different reaction of real estate loans across datasets, we include the 30-year mortgage rate in the credit model as a robustness check. We find that the response of real estate loans is indeed subject to whether the 30-year mortgage rate is included. Figure 3.5 in Appendix 3.B.2 illustrates that the point estimates of the IRFs of real estate loans are lower compared to the credit dataset without the 30-year mortgage rate. In line with the housing dataset, the effect of a monetary policy shock on real estate loans is now negative and significantly different from zero between about one and two and a half years after impact. Regarding M1 velocity and the M1/M2 ratio, we see that the two move in opposite directions, which is in line with what standard theory would suggest. Fluctuations in M1/M2 are driven by flows into and out of M1 (the non-interest paying component) and into and out of M2-M1 (the interest paying component). The reactions we observe for M1 velocity and M1/M2 differ greatly from the response to a permanent (or highly persistent) interest rate shock. In response to such a permanent shock, M1 velocity would behave similarly to the short-term interest rate, and M1/M2 would move in the opposite direction. However, in our case, the identified change in the interest rate is very short-lived. Hence, monetary aggregates follow a different pattern.

3.5.3 Costs of Leaning against the Wind

We now turn to the costs of leaning against housing prices or credit fluctuations. As our measure of costs, we consider the ratio between the cumulative IRF of either real housing prices or real credit and the cumulative IRF of real GDP. The costs are thus assessed in terms of real GDP losses. Table 3.1 reports the median and the 16-84 percentiles of the ratios at several horizons after the interest rate hike. We focus again on the results of the frequentist approach.¹⁹

Focusing on the median ratio, evidence suggests that the impact of an unexpected 25bp interest rate hike on real housing prices or real credit is not significantly larger than the impact on real GDP. Let us assume that the central bank would pursue a policy of leaning against housing prices. In this case, it would need to accept a cumulative fall in real GDP of 1% to achieve a cumulative decrease of housing prices by 4.1% three years ahead, 4.2% four years ahead, and 4% five years ahead. When it comes to real credit, the costs are higher. For a cumulative real GDP loss of 1%, real credit could be cumulatively decreased by only 1.6%

¹⁹The ratios based on the IRFs obtained from the Bayesian approach are provided in Appendix 3.B.3.

three years ahead, 1.7% four years ahead, and 1.5% five years ahead. Note that we consider non real estate loans as the relevant credit variable when we compute the ratio for the credit dataset.

Table 3.1: Costs of leaning against real housing prices or real credit in terms of real GDP losses

Years ahead:	3	4	5
Housing dataset	4.10 [3.37,6.37]	4.22 [2.97,7.44]	4.04 [-1.9,7.79]
Credit dataset	1.63 [1.2,1.79]	1.65 [1.1,1.84]	1.51 [0.64,1.79]

Note: Median ratios along with the 16-84 percentiles in brackets. We consider non real estate loans as the relevant credit variable for the ratio in the credit dataset. Results for the frequentist approach.

Regarding the confidence bands, we want to stress that the ones including negative values are hard to interpret and thus inconclusive. To illustrate this point, let us think about the two cases which can lead to such a negative ratio. First, it could be that the cumulative IRF of real housing prices is negative, while the cumulative IRF of real GDP is positive. In other words, a 25bp interest rate hike would decrease housing prices and, at the same time, increase real GDP. This case is very favorable for a central bank. Second, a negative ratio could stem from a positive cumulative IRF of real housing prices and a negative cumulative IRF of real GDP. In contrast to the first case, this second possibility is unfavorable for the central bank because housing prices increase further and real GDP declines after a monetary policy shock. Therefore, we focus mainly on the ratios, for which the lower as well as the upper band is positive. It is, however, reassuring that the median estimates for all considered horizons are very similar. This finding indicates that our median estimates for the costs of LAW policies are robust.

Evidence from the housing and credit dataset shows that by using its interest rate instrument, the ability of the central bank to keep ‘excessive’ growth in either real housing prices or real credit in check is very limited and associated with high real economic costs. For instance, if the central bank were to reduce the level of real credit by 10%, then it should be willing to accept a shortfall of real GDP of about 6%. The real GDP losses required to bring down real credit by a sufficient amount would thus be considerable. The same reasoning applies to housing prices. Thus, it is unlikely that a central bank would want to exploit such an unfavorable trade-off.

Overall, the evidence produced by random variation in the monetary policy rate around the path induced by the systematic component of monetary policy suggests that the costs of leaning against housing prices or credit fluctuations are substantial. Our results are in line with

the literature that advocates against LAW policies (Svensson (2017), Schularick et al. (2021), Benati (2021), Benati (2022)).

3.6 Conclusion

We evaluate central banks' costs when leaning against house prices or credit aggregates using their interest rate instrument. To measure these costs, we combine the state-of-the-art identification methods of Gertler and Karadi (2015) and Arias, Rubio-Ramírez, et al. (2021) with the significantly improved measures of MPS from Bauer and Swanson (2022) as an external instrument. We follow the current best practice in identifying monetary policy shocks for the US. The resulting IRFs to a monetary policy shock in the housing and credit datasets largely follow conventional wisdom. Based on these IRFs, we evaluate the costs of LAW policies by considering the ratio between the cumulative IRF of either real housing prices or real credit and real GDP at different horizons. The evidence produced by random variation in the monetary policy rate around the path induced by the systematic component of monetary policy suggests that LAW policies are very costly. The impact of an unexpected 25bp interest rate hike on real housing prices or real credit is not significantly larger than the impact on real GDP. On the one hand, if the central bank would follow a policy of leaning against housing prices, it would need to accept a real GDP loss of 1% to bring down real house prices by merely 4%. On the other hand, if it would want to lean against credit aggregates, a real GDP loss of 1% was necessary to achieve a decrease of roughly 1.5% in non real estate loans. The ability of the central bank to correct disequilibria in asset price markets or credit markets thus appears limited and is associated with significant real economic costs. Our results are in line with the literature that advocates against LAW policies (Svensson (2017), Schularick et al. (2021), Benati (2021), Benati (2022)).

3.A Appendix I

In the following, we provide further information on our dataset. Table 3.2 describes the data in detail and Table 3.3 explains how the data was transformed.

Table 3.2: Data Description

Raw Series	Mnemonic	Unit	Source
2Y Treasury Yield	GS2	%	FRED
Commodity Prices	PCOMM	Index	Bauer and Swanson (2022)
Personal Consumption Expenditures Excluding Food and Energy	PCEPILFE	Index	FRED
CPI for All Urban Consumers: Shelter in US City Average	CUSR0000SAH1	Index	FRED
Interpolated Real GDP	RGDP_INT	bn\$	Stock and Watson (2012) and IHS Markit
Interpolated Nominal GDP	NGDP_INT	bn\$	Stock and Watson (2012) and IHS Markit
Excess Bond Premium	EBP	%	Gilchrist and Zakrajšek (2012)
M1 Aggregate	M1SL	bn\$	FRED
Money Market Deposit Accounts	MMDA	bn\$	Benati et al. (2021)
M2 Aggregate	M2SL	bn\$	FRED
Real Estate Loans, All Commercial Banks	REALLN	bn\$	FRED
Consumer Loans, All Commercial Banks	CONSUMER	bn\$	FRED
Commercial and Industrial Loans, All Commercial Banks	BUSLOANS	bn\$	FRED
Other Loans and Leases: All Other Loans and Leases, All Commercial Banks	AOLACBW027SBOG	bn\$	FRED
Median Sales Prices of New Homes Sold in United States	NOMHP	\$	US Census Bureau
Monthly Supply of New Houses in the United States	MSACSR	thd	FRED
New Privately-Owned Housing Units Started: Total Units	HOUST	thd	FRED
All Employees, Construction	USCONS	thd	FRED
30-Year Fixed Rate Mortgage Average in the United States	MORTGAGE30US	%	FRED
Real Effective Exchange Rate, Narrow Indices	REALEXR	Index	BIS

Following Benati (2022), we use nominal quantities for several variables (commodity prices, real estate loans, house prices, and rent price index) when estimating our VARs. For the displayed IRFs and the calculation of the cost of LAW policies, we transform the IRFs into real responses. To be precise, we subtract the impulse response of the price index, i.e. the core PCE deflator, from the impulse response of said nominal quantities to get all IRFs in real terms.

Table 3.3: Data Transformation

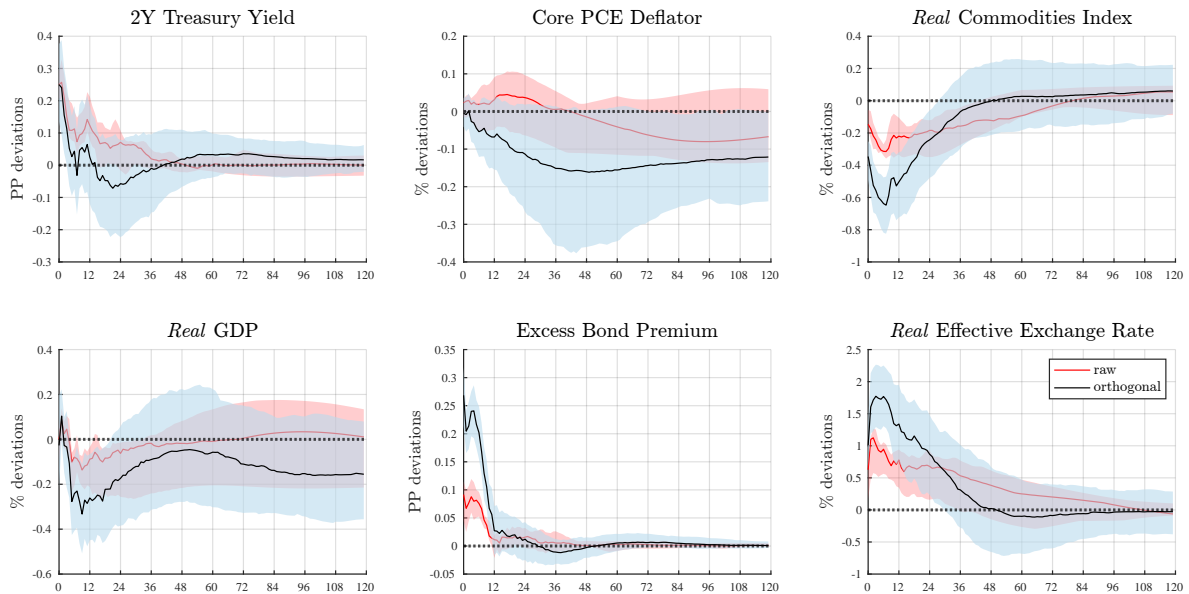
Constructed series	Formula
Two-Year Treasury Yield	$GS2$
EBP	EBP
Core PCE Deflator	$PCE = 100 \times \log(PCEPILFE)$
Nominal Commodities Index	$PCOMM = 100 \times \log(PCOMM)$
Real GDP	$RGDP_INT = 100 \times \log(RGDP_INT)$
Nominal Real Estate Loans	$REALLN = 100 \times \log(REALLN)$
Nominal Non Real Estate Loans	$NONREALLN = 100 \times \log(CONSUMER + BUSLOANS + OTHERLOANS)$
Nominal House Prices	$NOMHP = 100 \times \log(NOMHP)$
Nominal Rent Price Index	$RENTPRICE_INDEX = 100 \times \log(CUSR0000SAH1)$
M1 Velocity	$M1_VELO = NGDP_INT / (M1SL + MMDA)$
M1/M2	$M1/M2 = (M1SL + MMDA) / M2SL$
Employees Construction	$EMPL_CONSTR = 100 \times \log(USCONS)$
30Y Mortgage Rate	$MTG30Y = MORTGAGE30US$
Real Effective Exchange Rate	$REALEXR = 100 \times \log(REALEXR)$

3.B Appendix II

3.B.1 Instrument: Bayesian Approach

Below, we replicate Figure 3.1 of the main text using the Bayesian approach.

Figure 3.4: IRFs with raw versus orthogonalized surprises

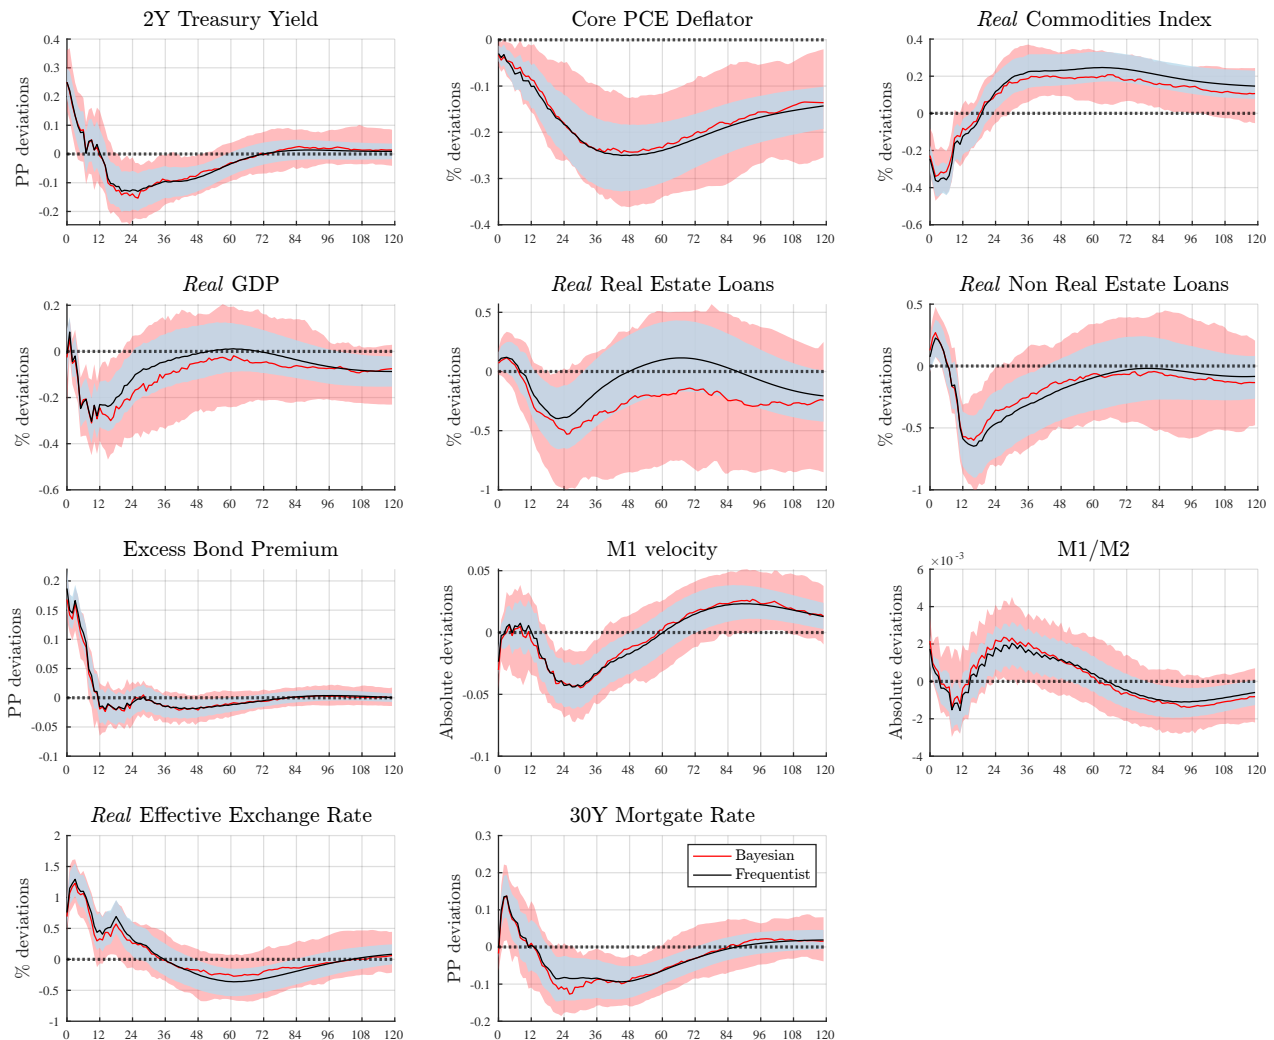


Note: The solid black line and the blue shaded areas are the point estimate and the 16-84 percentiles using the orthogonal MPS. The solid red line and the red shaded areas are obtained using the raw MPS. IRFs are normalized to a 25bp increase in the two-year Treasury yield. Horizontal axis: time horizon in months. Bayesian approach using 400,000 independent draws.

3.B.2 Evidence: Credit Model with 30Y Mortgage Rate

Below, we provide the IRFs for the credit model amended with the 30-year mortgage rate.

Figure 3.5: IRFs to a monetary policy shock in the credit dataset



Note: The solid black line and the blue shaded areas are the point estimate and the 16-84 percentiles resulting from the frequentist approach. The solid red line and the red shaded areas are the median and the 16-84 percentiles obtained with the Bayesian approach. IRFs are normalized to a 25bp increase in the two-year Treasury yield. Horizontal axis: time horizon in months.

3.B.3 Cost of Leaning against the Wind: Bayesian Approach

Below, we provide the cost of leaning against the wind calculated by using the Bayesian evidence.

Table 3.4: Costs of leaning against real housing prices or real credit in terms of real GDP losses

Years ahead:	3	4	5
Housing dataset	3.46	2.97	2.38
	[-3.12,6.98]	[-3.31,5.98]	[-3.12,6.34]
Credit dataset	1.21	1.12	1.03
	[0.15,2.35]	[-0.3,2.56]	[-0.95,2.62]

Note: Median ratios along with the 16-84 percentiles in brackets. We consider non real estate loans as the relevant credit variable for the ratio in the credit dataset. Results for the Bayesian approach.

Chapter 4

A Model United Nations Experiment on Climate Negotiations

Abstract

Weitzman (2014) proposed that focusing international climate negotiations on a uniform common commitment (such as a uniform carbon price) is more effective than negotiations on individual commitments (as in the Paris agreement) in achieving ambitious climate action. We put this hypothesis to an experimental test by simulating international negotiations on climate change in collaboration with Model United Nations (MUN) associations. This novel experimental format combines some of the advantages of lab and field experiments. Our results offer support for Weitzman's hypothesis and indicate that negotiating a common commitment on a uniform carbon price may yield higher emission reductions in the long run and more participation than individual commitments à la Paris.

Acknowledgment: Joint with Elisa Hofmann (Department of Intercultural Studies and Business Communications, University of Jena, Ernst-Abbe-Platz 8, Jena, Germany) and Klaus M. Schmidt (Department of Economics, University of Munich, Ludwigstraße 28, Munich, Germany). We would like to thank Friedrich Breyer, Uwe Cantner, Ernst Fehr, Michael Gerfin, Jonas Meier, Axel Ockenfels, Doina Radulescu, Tobias Regner, Asri Özgümüs, Wladislaw Mill, Ralph Winkler, the editor and two anonymous referees for helpful comments and suggestions. Financial support by Deutsche Forschungsgemeinschaft through CRC TRR 190 (project number 280092119) is gratefully acknowledged. Tanja Mitric, Timothy Rabozzi, Friederike Reichel, and Gabriel Vollert provided excellent research assistance. Moreover, we would like to thank the following MUN Associations for their support: Model United Nations University of Bern, ETH Model United Nations, MUN Team UZH, Model United Nations TU Munich e.V., MUNAM e.V., Model United Nations Mannheim e.V., Cologne MUN Society e.V., and United Nations Hochschulgruppe Tübingen. This chapter is published in the following article: Hofmann, Elisa, Kyriacou, Lucas and Schmidt, Klaus M.. 'A Model United Nations Experiment on Climate Negotiations' *Jahrbücher für Nationalökonomie und Statistik*, 2022. <https://doi.org/10.1515/jbnst-2021-0054>. The article is part of the special issue 'A New Decade of Research on the Economics of Climate Change: Towards an Integrated View on a Sustainable Use of the Biosphere' published in the *Journal of Economics and Statistics*.

4.1 Introduction

Mitigating CO₂ emissions is a global public good problem. To solve it, effective international cooperation is required that needs to be negotiated by sovereign countries. The success of these negotiations depends on how they are structured. In the negotiations on the Paris agreement, countries negotiated a common, non-binding goal by how much to limit global warming. This goal is to be implemented by ‘nationally determined contributions’, i.e. individual commitments by the participating countries of how much to contribute to the common good. This negotiation design was very successful in achieving maximum participation (all 197 member nations of the United Nations Framework Convention on Climate Change (UNFCCC) signed the agreement) and an ambitious common goal (limit global warming to less than 2° C), but it did not induce the parties to engage in sufficient climate action to achieve this goal. Because each country has to bear the full cost of its mitigation efforts alone while the benefits are distributed across all nations, it is widely feared that the Paris agreement will fail to sufficiently limit global warming.¹

There is a new proposal, advocated by Weitzman (2014), Weitzman (2017a), Nordhaus (2015), Nordhaus (2019), MacKay et al. (2015), and others, to structure climate negotiations in a radically different way. They argue that negotiations should focus on a uniform price for carbon emissions. This negotiation design strives for a uniform common commitment that builds on reciprocity. If a country pushes for a higher carbon price, it knows that this higher price will apply uniformly to all countries. Thus, both the benefits and the costs of this action are borne by all nations. This induces each country to strive for a carbon price that it believes is optimal for the world as a whole, rather than some climate action that it believes is optimal for itself given the mitigation efforts taken by the rest of the world. However, a possible drawback of this design is that some countries may prefer to stay out of the agreement and not to impose any carbon price at all.

Is this new approach, focusing on a uniform common commitment, likely to be more successful than negotiating individual commitments, as in the Paris approach? In this paper we report on the results of a novel type of field experiment that sheds light on this question. We collaborated with ‘Model United Nations’ (MUN) associations in Germany and Switzerland and simulated ‘Conferences of the Parties’ (COPs) of the United Nations (UN) with student delegates. The conferences followed the actual COP rules of the United Nations. Student delegates were supposed to represent the position of the country they were assigned to, they were provided with detailed information, they had to carefully prepare for the event, and they wrote a position paper for their country before the conference started. At each location we held two simultaneous COP conferences on climate action with ten countries each, representing

¹See e.g. Cléménçon (2016), Rogelj et al. (2016), Jiang et al. (2019), UNEP (2019).

all major regions and the main conflicting interests. Student delegates were assigned randomly to the two COPs and to the countries they represented. Both COPs had to pass resolutions specifying the reductions of carbon emissions in the years 2030, 2040, and 2050. In committee C1, these reductions were achieved by binding individual commitments to emission reductions of the participating countries in these years, combined with a non-binding common goal on how much to reduce worldwide emissions. In committee C2, the parties that passed the resolution agreed to a binding commitment on a uniform carbon price for each of these years. Thus, C1 negotiated individual commitments (as in Paris), while C2 focused the negotiations on a uniform common commitment (a uniform carbon price).

Our study has four main results: First, we find that countries achieve significantly and substantially higher reductions of carbon emissions in 2050 in C2 where they negotiate a uniform carbon price. Second, in C1 countries are equally ambitious in their non-binding common goals as in C2 (no significant difference), but the actual individual commitments in C1 do not live up to these goals. Third, to our surprise, significantly more countries participate in the resolution if a uniform common commitment is negotiated than if negotiations are focused on individual commitments. This is partly due to the behavior of countries like Russia and Saudi Arabia. They are opposed to a carbon price not because they suffer a lot if a high carbon price is introduced in their countries, but rather if other countries introduce a high carbon price which reduces demand for their fossil fuel exports. Thus, these countries try to convince other countries to keep the carbon price low in exchange for their participation in the overall agreement. Finally, there are substantial and significant differences in the reduction of emissions between the different countries in C1, while countries contribute more equally to the common good in C2. This is partly due to the uniformity of the carbon price, and partly due to the fact that in C2 more countries participate in the resolution.

Our paper is closely related and complementary to Schmidt and Ockenfels (2021) who conducted a laboratory experiment comparing the negotiation designs of Paris, Kyoto, and the new approach striving for a uniform carbon price. In the lab experiment, four subjects faced an asymmetric public good problem and negotiated a binding contract on their contributions to the public good. Negotiations took place through a computer network without personal interaction, the problem was framed in an abstract and neutral fashion (climate change is never mentioned), and subjects were paid for their decisions. Schmidt and Ockenfels (2021) find that negotiations on a uniform minimum contribution to the public good are significantly and substantially more effective than a negotiation design with individual commitments as in Paris and negotiations on a common complex commitment as in Kyoto. This result is driven by two effects. First, negotiating a uniform minimum contribution to the public good induces all parties who participate in the negotiations to contribute almost efficiently (as predicted by the game-theoretic analysis). Second, while the participation rate is somewhat lower when a

carbon price is negotiated as compared to Paris style negotiations, the free-riding effect is small. Furthermore, Schmidt and Ockenfels (2021) show that their qualitative results continue to hold even if contracts are non-binding and cannot be enforced.

The advantage of the laboratory experiment of Schmidt and Ockenfels (2021) is its internal validity. The lab allows for a tight control of the environment and for many independent observations. Thus, the experiment can show that negotiation design has a statistically highly significant (and substantial) causal effect on the negotiation outcome. However, the experiment is very stylized, negotiations take place anonymously via a computer network, and subjects do not negotiate on climate action but on monetary outcomes of an abstract public good game. Thus, it is difficult to assess whether the experimental results of the lab carry over to the real world.² Our study addresses some of these problems. We look at negotiations following the rules of real COP negotiations, negotiations have a richer and more realistic context, they last for an entire day, and subjects are intrinsically motivated to represent the best interest of their countries. Following the taxonomy of Harrison and List (2004), our experiment is close to what they call a framed field experiment, because the experimental framework of climate change negotiations represents field context. However, our design differs from a framed field experiment as we use a standard subject pool of students. These subjects simulate negotiations in a field setting, an experimental design that is not quite captured in the taxonomy of Harrison and List (2004).

Our paper is related to several strands of the literature. First, there is a large literature on international environmental agreements going back to the canonical paper by Barrett (1994) and surveyed by Barrett (2005). This literature offers many important insights in the incentives to join international agreements and how to make them self-enforcing. However, it does not discuss the design of the negotiation process.

Second, there is a small literature on how to structure climate negotiations. In a series of papers, Weitzman (2014), Weitzman (2015), Weitzman (2017a), Weitzman (2017b) compares negotiations on a uniform carbon price to negotiations on a vector of emission reductions. He argues that negotiating a uniform carbon price provides a salient focal point (as advocated by Schelling (1960)), and he shows formally that it aligns self-interest with the common good. This argument has been formalized by Weitzman (2014) and Schmidt and Ockenfels (2021). They model the fight against climate change as an asymmetric public good problem. All countries want to mitigate global warming, but some countries are more affected by climate change or have lower costs of introducing a carbon price than others and therefore want to implement more stringent policies. Thus, for each country, there is a most preferred carbon price that it wants to be applied to the world. Schmidt and Ockenfels (2021) show that this mechanism is strategy-proof, so for each country it is a weakly dominant strategy to propose the carbon price

²An interesting discussion of external validity is given by List (2020).

that would be optimal from its own perspective, if it was imposed by all other countries as well. MacKay et al. (2015) furthermore emphasize that a uniform carbon price is a reciprocal instrument ('I will if you will'). Nordhaus (2015), Nordhaus (2019) proposes a 'climate club' that commits to a uniform carbon price for its members and imposes tariffs on the imports of non-members to compensate for the distortion of competition and to induce other countries to join the club. He points to the experience with other international agreements (Barrett (2003), Battaglini and Harstad (2016), Battaglini and Harstad (2020) showing that treaties tend to be stable only if they penalize free-riders.

Third, there is an experimental literature on the endogenous formation of institutions for successful public good provision. Dannenberg and Gallier (2020) offer an excellent survey of this literature. According to their classification our paper considers the case of a global public good (carbon emissions affect everybody) and an exclusive institution (the uniform carbon price applies to only those countries that support it). A closely related paper in this class is Kosfeld et al. (2009) who consider a public good game in which subjects can vote on the introduction of an institution that forces all members to contribute the full amount to the public good, while those who stay out are free how much to contribute. They find that institutions form in about half of all cases, and in most of them there is full participation. In contrast, in our setup countries are not forced to raise the carbon price to the efficient level. The minimum mechanism allows them to react to those countries that did not participate in the uniform carbon price. Several other papers explore the minimum contribution mechanism. Dannenberg, Lange, et al. (2014) consider a repeated public good experiment in which all players are forced to participate (inclusive institution). They find that about 60 percent of the groups have an increasing minimum contribution level over time that approaches the social optimum at the end, while the other 40 percent implement a low level throughout the game. This may be explained by the fact that participation was not voluntary. A subject that was forced to participate can choose a very low contribution level which implies that all other contributions are equally low. In contrast, participation was voluntary in our setup. In Kocher et al. (2016) as well as in Martinsson and Persson (2019) the minimum contribution level was prespecified but lower than the efficient level. Both papers find that the large majority of subjects vote in favor of this mechanism. Almost all of the lab experiments in this literature do not allow subjects to communicate, while in our Model United Nations experiment parties had ample time to discuss and negotiate their actions. While almost all of the lab experiments take place in an abstract setting, there is one interesting experiment by Barrett and Dannenberg (2016) on climate negotiations. They study the 'pledge and review' process of the Paris agreement in a lab experiment and show that pledges do not increase actual contributions over time.

Finally, there is a literature in political science and education on simulation games (see, e.g., Boardman (1969), Lester and Stoil (1979), Asal (2005), Kauneckis and Auer (2013). This

literature focuses on simulation games mainly as a pedagogical tool to foster student engagement and learning. A few exceptions are Penetrante (2012) and Matzner and Herrenbrück (2017). Penetrante (2012) used MUN simulation games as ‘case studies’ to find stumbling blocks in negotiations and to analyze coalition building. Similarly, Matzner and Herrenbrück (2017) conducted three MUN conferences as experimental sessions to explore the conflicts that may arise when countries negotiate on climate engineering. However, none of these papers allows for statistical hypothesis testing. A methodological novel feature of our paper is the use of MUN simulations as a randomized controlled field experiment. This is related to Schwardmann et al. (2022) who used international debating competitions as a ‘field experiment’ to study whether people persuade themselves about the moral and factual superiority of their position to better convince others.

The rest of the paper is structured as follows: Section 4.2 describes the setup of the MUN experiment. Section 4.3 derives the hypotheses that we want to test. Section 4.4 summarizes the experimental procedures and Section 4.5 reports the results. Section 4.6 concludes with a discussion of our new experimental method and our main results. A detailed documentation of the MUN experiment is provided in the Online Appendix.³

4.2 The MUN Experiment: Design

For the experiment we collaborated with Model United Nations (MUN) associations at six universities across Germany and Switzerland. The idea of MUN goes back to the League of Nations simulations in Oxford and Harvard in the 1920s (Muldoon (1995)). Today, MUN associations exist at hundreds of universities and high schools in most countries of the world. They are student organized and engage in extracurricular activities to prepare for simulated United Nations conferences. They teach their members debating and writing skills as well as critical thinking and leadership abilities, they organize local MUN events and send delegations to national and international MUN conferences.

A MUN conference simulates real United Nations conferences. Participants in a MUN conference (‘delegates’) are assigned countries to represent. They have to conduct research on their country, formulate positions and come up with policy proposals that they will then debate with their fellow delegates at the conference, staying true to the actual position of the member they represent. During the conference, delegates have to adhere to the formal rules similar to those of real UN conferences. At the end of the conference delegates vote on written policies, called ‘resolutions’ with the goal of passing them with a majority vote. The best performing delegates are often recognized with awards.

³The Online Appendix can be downloaded from degruyter.com.

At each of the six universities we organized a one-day MUN conference on climate policy⁴ with two separate committees (treatments) consisting of ten delegates and two chairs each. Delegates and chairs are allocated randomly to the two committees and to the countries they represent. There are ten nations, kept constant across all conferences, that are supposed to represent the world: Australia, Canada, China, European Union (EU), India, Japan, Russia, Saudi Arabia, South Africa, and the United States of America (USA).⁵ These include the major CO₂ emitters (USA, China, EU, Japan), the major producers of fossil fuels (Saudi Arabia, Russia, USA, Australia), and countries from the developing world (India, South Africa). The two chairs are supposed to strictly enforce the rules of the conference, but not to influence the negotiations otherwise.

Two weeks before the conference all participants receive the same ‘study guide’ (available in the Online Appendix) which offers general information on the causes and consequences of climate change and a summary of the forecasted scenarios of the 5th Assessment Report of the Intergovernmental Panel on Climate Change (IPCC). It also discusses different policy options to combat climate change and their estimated costs.⁶ Participants are supposed to research the position of the country they represent and to write and hand in a position paper before the conference starts.

The position paper has to include (1) a statement on how the represented country will be affected by climate change, (2) what the country’s general position on climate action is, (3) a suggestion by how much the world as a whole should reduce CO₂ emissions until 2030, 2040, and 2050 (as compared to 2010), (4) a suggestion by how much the respective country should

⁴The MUN conference is modelled after a ‘Conference of the Parties’ (COP). The COP is the main body of the United Nations Framework Convention on Climate Change (UNFCCC), which represents the framework of United Nations climate change activities. At the COP, all United Nations Member States meet once a year to negotiate and adopt international climate change agreements. For instance, the Kyoto protocol has been negotiated upon at COP3 in 1997, and the Paris agreement has been adopted at COP21 in 2015. Thus, COP is a natural candidate being simulated at a MUN conference when investigating international climate change negotiations.

⁵We are aware that the European Union is not a country but represents all member countries of the European Union. This was explicitly stated in the study guide. We nevertheless included the European Union as ‘one country’ because the European countries closely coordinate their climate policies at the European level.

⁶The study guide also offers some statistical information on the ten participating countries, such as population, GDP/person, share of worldwide CO₂ emissions, and emissions per person in 2010, etc. Finally, it provides references and links to research further information. In addition, participants receive the ‘MUN Rules of Procedures’, which setup the formal rules for the debate (available in the Online Appendix). Furthermore, participants are provided with a document entitled ‘Model United Nations for Beginners’. It reminds the delegates about the commonly used documents during MUN conferences (such as Working Papers, Draft Resolutions, and Resolutions), a short introduction to MUN Rules of Procedures (e.g. rules on debating, on speeches, and on diplomatic conduct) and a summary of voting rules. Finally, delegates are provided with an Excel Sheet that could be used to compute the effects of different policy measures in different countries on worldwide CO₂ emissions. We selected the best position paper based on whether it took a credible and realistic stance on the aforementioned five points from the perspective of the represented country and how well written the position paper was in terms of structure, information content and clarity. These award-winning position papers (from both committees and all experimental sessions) are available in the Online Appendix.

reduce its own CO₂ emissions until 2030, 2040, and 2050 (as compared to 2010), and (5) a short discussion of the advantages and disadvantages of carbon pricing from the perspective of the represented country. We incentivize the delegates to stay in character of their nation by publicly awarding a ‘Position Paper Award’ (cinema voucher of about 10 Euros) to the best written and most authentic position paper at the end of the conference.⁷ Furthermore, we proofread position papers and remind delegates whose position papers do not answer all five questions to think about the unanswered questions. The position papers are distributed among the delegates and chairs in each Committee. All participants are asked not to communicate about the conference before the conference starts.

Before the conference, the two committees receive a document called ‘Questions A Resolution Must Answer’ (QARMA, available in the Online Appendix). This document is the only difference between the two committees. It sets up different negotiation frameworks. Committee 1 (C1) is instructed to get to a resolution that specifies goals for the reduction of worldwide CO₂ emissions in 2030, 2040, and 2050, but these goals are non-binding.⁸ Actual reductions are achieved by nationally determined contributions, i.e. individual commitments of all countries supporting the resolution on how much they will reduce their own emissions by 2030, 2040, and 2050. Parties are also asked to indicate in the Annex how they plan to achieve these reductions (e.g. via carbon pricing, subsidizing green energy, imposing restrictions and regulations on CO₂ emissions, etc.) Participants are told that these country-specific goals are binding commitments according to international law.⁹ Note that the country-specific goals are ‘nationally determined’ by each country alone after the resolution has been passed. This negotiation protocol shares key features of the negotiation protocol that led to the Paris agreement.

In Committee 2 (C2) the negotiation rules are different. Delegates are informed that there has been a previous agreement to introduce a common commitment via a uniform price for CO₂ emissions (a ‘carbon price’). They are instructed to get to a resolution that specifies how high this uniform carbon price should be in 2030, 2040, and 2050. Each country is free how to implement the carbon price (by a carbon tax, an emissions trading system, or some hybrid system) and how to spend the revenues from carbon pricing. However, for all countries supporting the resolution, the carbon price is binding. Again, participants are told that this agreement is a binding commitment according to international law. Countries that do not

⁷We selected the best position paper based on whether it took a credible and realistic stance on the aforementioned five points from the perspective of the represented country and how well written the position paper was in terms of structure, information content and clarity. These award-winning position papers (from both committees and all experimental sessions) are available in the Online Appendix.

⁸Reductions or increases in CO₂ emissions are always relative to the emission levels of 2010.

⁹The instructions mention that if a country does not adhere to its legally binding commitment, it can be sued and sentenced to pay penalties. Note that the legal status of the actual Paris agreement is less clear. While some countries have been successfully sued by their citizens to fulfill their obligations, this is not possible in other countries. However, we wanted to have equally strong commitments in both treatments. We did not discuss whether these legal obligations are enforceable.

vote for the resolution are not bound by it. Delegates are informed that the common carbon price is accompanied by a carbon border adjustment tax to prevent carbon leakage.¹⁰ They are instructed that the carbon price has to be uniform (i.e. the same for all countries supporting the resolution) and that the resolution cannot include any additional measures to mitigate climate change.

In both committees, a resolution needs at least five votes to be passed.¹¹ Countries that do not support the resolution are assumed to pursue ‘business as usual’ which results in an increase of CO₂ emissions by 40 percent of this country. Furthermore, in both committees, delegates are instructed not to discuss compensation payments for low income countries because there will be another UNFCCC conference on the introduction of a ‘Green Fund’ that will deal with this issue separately.¹²

It is crucial for the experiment that the climate effects of the negotiation outcomes can be compared with each other. In the study guide we provide a table that informs delegates that there is a linear relationship between the reduction of CO₂ emissions and the carbon price. For example, a carbon price of USD 0 yields an increase of emissions of 40 percent, a carbon price of USD 60 yields an increase of emissions of 0 percent, a carbon price of USD 120 yields a reduction of emissions of 40 percent and a carbon price of USD 180 yields a reduction of emissions of 80 percent. This linear relationship applies to all countries and to the world as a whole. It is clearly a simplification, but it corresponds roughly to estimates that have been reported by the IPCC AR5 (IPCC (2014)) and the Report of the High-Level Commission on Carbon Prices (Stiglitz et al. (2017)).¹³ All delegates are instructed to take these numbers at face value.

The two committees are two experimental treatments that differ in only one respect: In

¹⁰The idea of a climate club with a uniform carbon price is closely linked to a carbon border adjustment mechanism, see e.g. Weitzman (2014, p. 47) and Nordhaus (2019, p. 2011). For this reason, we presented these two ideas together in Committee 2. A carbon border adjustment mechanism is much more problematic if employed by a single country. Hence, we did not refer to it in the instructions of Committee 1.

¹¹This voting procedure is less demanding than the one applied at the COP21 in which the Paris Agreement has been negotiated (which required a two-third majority and for some decisions a consensus, see UNFCCC Secretariat (2015)). Out of 12 voting results from both committees, we observed 4 instances where the resolution was adopted unanimously, 7 instances where it was adopted with a two-thirds majority and one instance where it was adopted with a simple majority (in this case the resolution passed with 6 Yes votes). On average, 8 out of 10 countries voted in favor of the resolution in C1. In C2 the average was 9.3. See the discussion of Result 3 and Figure 4.4 for more details.

¹²We added this statement to the instructions after the first conference in Bern. In Bern delegates spent some time discussing this issue and decided in the end not to deal with it (in both committees). We wanted to focus the attention of the delegates on the measures against climate change, so we added this statement. It does not seem that the resolutions were affected by these discussions. All our results remain valid if we exclude Bern from the analysis (see Online Appendix).

¹³The linear relationship reflects the short- and medium-term effects of a carbon price. The long-term effects are predicted to be significantly larger because of technological innovation. For example, the cost of producing green hydrogen and the cost of carbon capture and storage are predicted to decrease significantly over time and will become economical at much lower carbon prices. If this prediction is correct, it strengthens the effects of carbon pricing (and thereby our results) significantly.

Committee C1, individual commitments are negotiated, while in Committee C2 negotiations aim at a uniform common commitment. Note that carbon pricing is not restricted to Committee C2. The delegates in Committee C1 can also choose to adopt carbon prices. However, in C1 each country has to do this independently, while in C2 all countries supporting the resolution have to commit to a common uniform carbon price.

4.3 Hypotheses

The main question of our study is which negotiation protocol is more successful in achieving an effective agreement.¹⁴ Following the literature discussed in the Introduction we predict that negotiations on a uniform carbon price are more successful than negotiations on individual commitments.

Hypothesis 1: Negotiations in C2 achieve significantly higher reductions in CO2 emissions than negotiations in C1.

In C1, parties did not only state their nationally determined contributions, they also proclaimed a non-binding common goal, namely by how much they want to reduce global CO2 emissions. Because this common goal is non-binding and nobody can be held responsible for not achieving it, we hypothesize that the individual commitments fail to achieve the stated common goal.

Hypothesis 2: In C1 the proclaimed non-binding common goal is significantly higher than the sum of the nationally determined contributions specified by each country individually.

How many countries will participate in signing the resolution? The Paris agreement was signed by all 197 nations on the planet, possibly because each nation was free how much to contribute to the commonly declared goal. But universal participation is not guaranteed. The USA was going to withdraw from the agreement while our experiments took place and some other countries were also on the fence. Thus, in C1 we expect most (but not necessarily all) countries to participate. In C2, a uniform carbon price is a much stronger commitment. Each party has an incentive not to participate but to free-ride on the efforts of the other parties.

Hypothesis 3: Participation in the resolution is significantly higher in C1 than in C2.

A related question concerns the differences across countries in their contributions. In C2 the uniform carbon price forces all countries that participate in the resolution to impose the same

¹⁴We (pre-)registered the trial with a Pre-Analysis Plan in the AEA RCT Registry: Hofmann, Elisa, Lucas Kyriacou and Klaus Schmidt. 2019. 'A Model United Nations experiment on climate change negotiations.' AEA RCT Registry. October 23. <https://doi.org/10.1257/rct.4834-1.1>. However, we did so only after the first sessions had been completed.

price. Only those countries that do not participate have a carbon price of 0 and therefore higher emissions. In contrast, in C1 each country decides on its own how much to contribute. Some countries (like the EU or Japan) may behave altruistically, others (like the USA or Russia) may put their own interests first. Therefore, we hypothesize that the differences in contributions are likely to be larger in C1.

Hypothesis 4: In C1 there will be large differences in the contributions of the different countries. In C2, emission reductions will be less unequal.

4.4 Experimental Procedures

The experiment was conducted in late 2018 and 2019. We held six MUN conferences in Germany and Switzerland lasting about seven hours (10:00 a.m. until 5:00 p.m.) each. Altogether we collected data from 144 participants (120 delegates and 24 chairs, see Table 4.1).¹⁵ The

Table 4.1: Overview of experimental sessions and participants.

Session Date	Session Location	Participants		
		Committee 1	Committee 2	Sum
15.12.2018	Bern (Switzerland)	12	12	24
23.03.2019	Munich (Germany)	12	12	24
18.05.2019	Zurich (Switzerland)	12	12	24
19.10.2019	Mannheim (Germany)	12	12	24
09.11.2019	Cologne (Germany)	12	12	24
30.11.2019	Tübingen (Germany)	12	12	24
	<i>Total</i>	<i>72</i>	<i>72</i>	<i>144</i>

communication language of the conference was English. Subjects received a flat payment of 20€ in Germany and 20CHF in Switzerland for their participation. Participants knew that the conference was part of an experiment, but they did not know what the experiment was about, nor did they know that the other committee negotiated under a different negotiation protocol. We separated the two committees from the very beginning and instructed the participants not to interact with members of the other committee before and during the conference. After a short introduction given in each committee by an experimenter, the conference was run by the chairs. One experimenter was sitting in the back of each room and kept a log.

Table 4.2 provides summary statistics of the participants (delegates only)¹⁶ and survey evidence on how they evaluated the conference. On average, participants were 22.07 years old.

¹⁵In our statistical analysis we ignore the chairs and are interested in the behavior of the delegates only ($n = 120$).

¹⁶Missing observations are due to not obligatory responses in the ex-post questionnaire.

48 percent of them were female. Subjects studied on average in their 4th semester. On average participants had participated in 3.25 MUN conferences before and prepared on average 467.22 minutes (almost eight hours) for the conference. They rated the realism of their position paper at 5.32 on a seven-point Likert-Scale ranging from 1 ('not realistic at all') to 7 ('very realistic'). They evaluated their own performance at 4.82 (on the same scale), similar to the performance of the other participants (4.92). Furthermore, subjects evaluated the realism of the resolution at 4.24 and they were 'rather satisfied' with the resolution (4.71 on a seven-point-Likert-Scale ranging from 1 ('not satisfied at all') to 7 ('very satisfied')). The participants stated being very interested in the topic already before the conference but indicated an even (significantly) higher interest in the topic after the conference (one-sided Wilcoxon signed-rank test, $z = 4.55$, $p < .001$).¹⁷ This survey evidence indicates that the MUN conferences succeeded in providing a fairly realistic framework for our simulation experiment. Participants came from a broad

Table 4.2: Summary descriptive statistics for delegates across all six experimental sessions.

Variable	Mean	Std. Dev.	N	Min.	Max.
Age	22.07	2.74	119	16	29
Gender (female = 1)	0.48	0.50	120	0	1
Semester	4.33	2.64	98	1	12
MUN participations	3.25	3.14	116	0	20
Preparation time (in minutes)	467.22	330.78	116	50	2040
Evaluation: Realism position paper	5.32	1.25	119	1	7
Evaluation: Realism own performance	4.82	1.21	119	1	7
Evaluation: Realism other's performance	4.92	1.14	119	2	7
Evaluation: Realism resolution	4.24	1.45	119	1	7
Satisfaction with resolution	4.71	1.63	119	1	7
Interest in topic before conference	5.55	1.27	119	1	7
Interest in topic after conference	5.98	1.13	119	2	7


variety of disciplines. Most of them studied economics (28%), law (22%) and other social sciences (21%). The Balance Table A.1 in the Online Appendix shows that the differences among participants across treatments and experimental session are small.

¹⁷This result confirms previous studies that have shown that simulations tend to increase the participants' interest in the topic, see e.g. Boardman (1969) and Rooney-Varga et al. (2018). In our study we mainly rely our statistical analysis on nonparametric tests (Wilcoxon signed-rank test, Kruskal-Wallis test, Fisher's exact test) due to our small sample size and the potentially non-normality of distributions. We complement these nonparametric analyses by parametric analyses (linear mixed effects regressions for Result 1 and binary logistic regression for Result 3).

4.5 Results

Figure 4.1 compares the average actual worldwide reductions of emissions in percent as compared to 2010 in C1 (individual commitments) and C2 (uniform common commitment) across the six COP meetings.¹⁸ Note that positive values indicate emission reductions while negative values indicate increases in emissions.

Figure 4.1: Actual worldwide reductions of emissions (weighted average over all countries) in Committee C1 versus Committee C2 averaged over all COP meetings.



There are two main differences between the outcomes of treatments C1 and C2. First, average emission reductions of the years 2030-2050 are substantially higher in C2 (about 20 percent) than in C1 (about 5 percent). Second, the time path of emission reductions in C2 is significantly steeper than in C1. Both start with a similar reduction of emissions in 2030 of about -8 percent (i.e. an increase of emissions by +8 percent). However, in 2040 and in particular in 2050, the negotiation outcome in C2 is much more ambitious as compared to C1. In 2050 this difference is statistically significant.¹⁹

Result 1: In 2050 the actual emission reductions of the ten countries are significantly higher in C2 than in C1. Furthermore, the time path of actual average reductions is significantly steeper in C2 than in C1.

¹⁸The emission reductions for each experimental session are provided in Table A.2 in the Online Appendix.

¹⁹In fact, in C2, the average worldwide reduction of emissions is about 50 percent in 2050 while it is less than 20 percent in C1. Thus, if the experiment and the underlying assumptions are taken literally, then the reduction of emissions through a uniform minimum price of carbon would just be sufficient to reach the two-degree-goal, while the reductions achieved via Paris style negotiations would be far too small to achieve this.

Result 1 confirms Hypothesis 1 in 2050. It is supported by a one-sided Wilcoxon signed-rank test ($z = 2.071$, $p = .019$) comparing the actual CO2 emission reductions of the ten countries in Committee C1 (as stated in the Annex if they support the resolution and ‘business as usual’ otherwise) with the actual CO2 emission reductions of the ten countries in Committee C2 (again depending on voting behavior on the resolution) for the year 2050 in all experimental sessions.²⁰

Result 1 is further backed up by linear mixed effects regressions (see Table 4.3).

Table 4.3: Determinants of actual average CO2 emission reductions 2030-2050 grouped on experimental session level.

	(1)	(2)	(3)
(Intercept)	5.54 (6.46)	-36.41*** (9.03)	-19.29** (6.36)
C2 Treatment	14.48 (10.39)	14.48 (10.39)	-19.76 (12.02)
Year		20.98*** (3.75)	12.42*** (2.20)
C2 Treatment x Year			17.12* (9.79)
AIC	355.08	336.97	333.93
BIC	363.00	344.89	343.43
Log Likelihood	-172.54	-163.49	-160.97
Num. obs.	36	36	36

Notes: *** $p < .01$, ** $p < .05$, * $p < .1$; Regression results from mixed effects models; Abatement levels of carbon emissions was the dependent variable in all three Models; Random intercepts are associated with location in all three Models and additionally with time in Model 1. Robust standard errors adjusted for six location clusters are provided in parentheses.

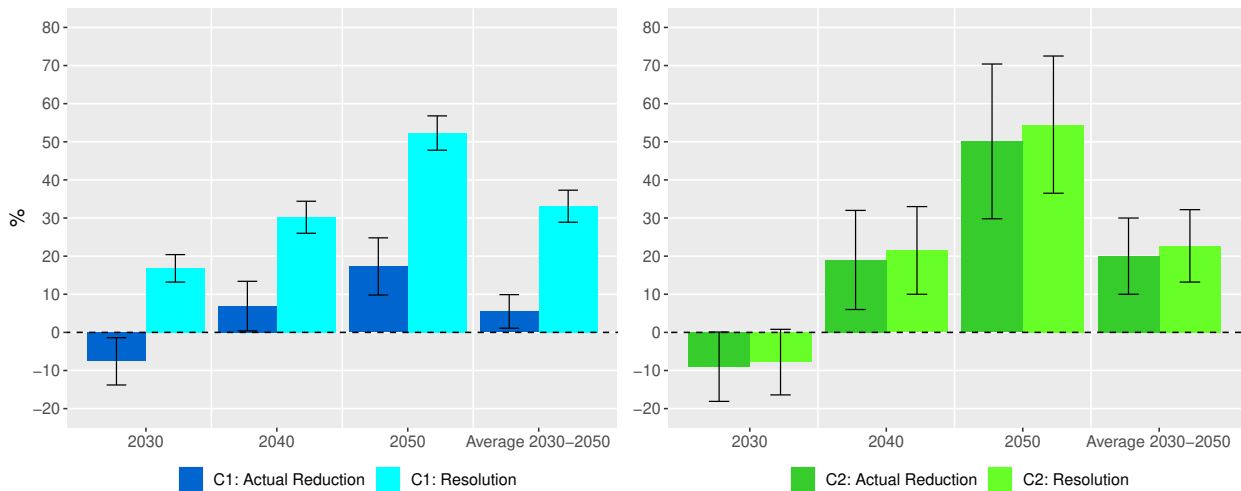
Here we compare the actual average worldwide abatement levels over 2030-2050 between the two treatments and for the different years. In all three models, the dependent variable is the actual average CO2 emissions reduction in the years 2030, 2040 and 2050. We include a random effect of the location as a matching group variable in all three Models. In Model 1, we further add a random effect of time to account for possible dependencies of the three responses (2030, 2040, and 2050) of one committee (repeated observations) and hence to account for the

²⁰We use a one-sided test here (and in some of the following results) to test directed predictions, while two-sided tests are used for undirected predictions.

nested structure of our data.²¹ The treatment itself increases actual worldwide reductions of emissions, but the effect is statistically not significant (Model 1, coefficient = 14.48, $p = .163$). However, reductions of emissions significantly increase over time (Model 2, coefficient = 20.98, $p < .001$). Furthermore, we find a (marginally) significant interaction between treatment and time (Model 3, coefficient = 17.12, $p = .080$).²²

What drives the difference in emission reductions between the two treatments? If we look at the resolutions that are passed by the committees, then the negotiations in C1 are on average more ambitious than the C2 resolutions, in particular in 2030. However, there is a large and statistically significant difference between the resolution and the actual reduction of emissions in treatment C1, but not in treatment C2. This is shown in Figure 4.2.

Figure 4.2: Resolutions vs. actual reductions. The left bars depict the reductions of emissions as announced in the resolutions, the right bars the actual reductions of emissions in all countries (national commitments weighted by country size).



Result 2: There is no significant difference in the resolutions passed in C1 and C2. However, the emission reductions in the resolutions are substantial and statistically significantly higher than the actual reductions of emissions for 2030, 2040, and 2050 in C1, but small and not significantly different in C2.


Result 2 confirms Hypothesis 2. It is supported by Wilcoxon signed-rank tests (two-sided test, $z = -1.614$, $p = .107$), indicating that the resolutions passed are not statistically significantly

²¹We ran experimental sessions at six locations with two treatments each, resulting in 12 group observations being matched at the location level. Due to the three (repeated) responses per observation (2030, 2040, and 2050), we include 36 observations in our regression analysis. As a robustness check we ran a similar regression including a random effect of the group (clustering the standard errors at the 12 group levels; see Online Appendix, Table A.3). When excluding the data from Bern, the regression results are similar (see Online Appendix).

²²As a robustness check we ran a similar regression with time dummy variables (see Online Appendix, Table A.4) confirming these results.

different between C1 and C2. Furthermore, the emission reductions as announced in the resolutions are statistically significantly higher than the actual average emission reductions for 2030, 2040, and 2050 in C1 (one-sided tests, $z = 2.201$, $p = .014$), while they do not differ in C2 (two-sided tests, $z = 1.408$, $p = .159$). There are two reasons for Result 2. First, in C2 the minimum price for carbon is binding for all countries that passed the resolution. In contrast, in the C1 negotiations, the actual reductions of emissions are determined by nationally determined contributions while the non-binding goal of worldwide emission reductions is just cheap talk. In fact, while some countries reduce their emissions by more than the goal of the resolution, most countries reduce substantially less and do not live up to the proclaimed goal in the resolution that they passed. This is illustrated in Figure 4.3 that considers only negotiations in C1 and looks only at those cases where a country voted in favor of a resolution. As can be seen from the figure, while the EU, Japan, Canada and Australia reduced their emissions by more than required by the resolutions, India, South Africa, Russia, and China reduced their emissions much less. The USA did not reduce them at all, and Saudi Arabia even increased their emissions on average in the years 2030-2050.

Figure 4.3: Proclaimed goal vs. actual reductions of each country in C1. The left bars depict the average announced reduction in those resolutions that the country voted for, the right bars depict the average actual reductions in all cases in which the country voted for the resolution.

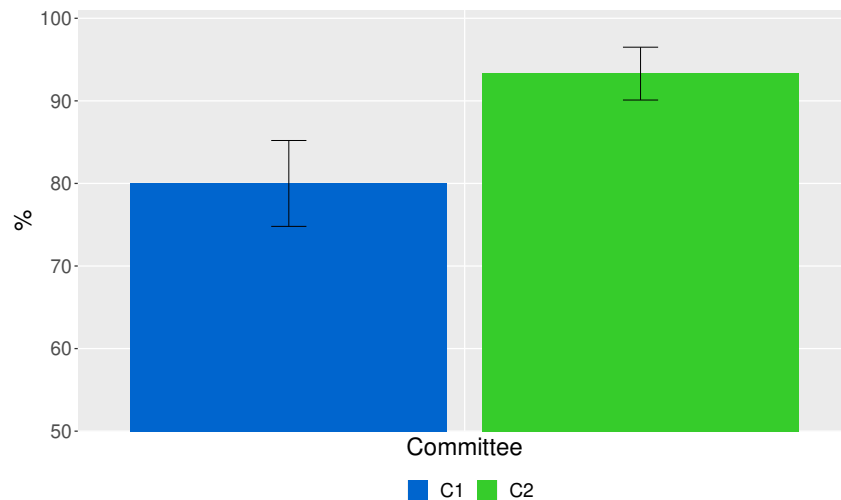


The second reason is that resolutions are supported by fewer countries in C1 negotiations than in C2 negotiations. This is shown in Figure 4.4. While in C1 the resolution was passed on average by 80 percent of all countries, it was accepted by 93 percent in C2.

Result 3: In C2 significantly more countries support the resolution than in C1.

Result 3 is supported by a Fisher’s exact test, which shows that a nation is statistically significantly more likely to vote in favor of the resolution if it is in C2 than if it is in C1 ($p = .029$). We further run a binary logistic regression, including the voting behavior as dependent variable and the treatment as predictor variable. Being in C2 significantly increases the probability of voting Yes versus voting No (coefficient = 1.25, $p = .041$, odds ratio = 3.5).

Figure 4.4: Fraction of countries that voted in favor of the resolution in Committee C1 and Committee C2 (averaged over all COP meetings).



Result 3 clearly refutes Hypothesis 3. This may be surprising at first glance. After all, in C1 each country is free to choose its nationally determined contribution. Thus, each country could simply vote for the resolution and then choose a much smaller reduction of its emissions than required by the resolution. In fact, this is what many countries did. However, some countries voted against the resolution, in particular the US (5 times), Russia (3 times) and Saudi Arabia (3 times), often because they wanted to make the point that the fight against climate change is harmful to their national interests. This reflects the fact that the Paris agreement was seen more critically in these countries at the time of the experiments than in 2015 when the Paris agreement was signed.

On the other hand, in C2, the agreed upon common carbon price is binding for all countries that supported the resolution. Thus, a country can free-ride only by rejecting the resolution. This suggests a lower acceptance rate in C2 than in C1. However, in the discussions and actual negotiations during the MUN conferences we frequently observed that countries opposed to carbon pricing used their participation in the resolution as a bargaining chip. The delegates of Russia and Saudi Arabia (and less frequently of the USA) argued that a high carbon price of the other countries (combined with a border adjustment tax) is a major threat to their interests as exporters of fossil fuels. At the same time, they recognized that the majority of nations is determined to introduce carbon pricing to mitigate climate change. Thus, they tried to keep

the carbon price as low as possible by leveraging their vote for the overall resolution. They promised to vote for the resolution, if the other countries agreed to carbon prices that were not too high.²³ This strategy often proved successful. Our observation of the actual negotiations indicate that the other countries probably would have adopted substantially higher carbon prices, in particular in 2030, if countries like Russia, Saudi Arabia and the USA had stayed out.

The developing countries, in particular India, South Africa, and to a lesser degree China, sometimes objected to the idea that they should impose the same carbon price as the developed countries. Two main arguments convinced them to participate. First, they realized that they will suffer most from climate change and thus have a strong interest to induce the other countries to mitigate it. Thus, they leveraged their vote by pressing for a higher carbon price. Second, the prospect of the Green Fund played a role in all negotiations (both C1 and C2), even though this was not part of the official agenda. Developing countries speculated that they will succeed in a later conference to convince richer countries to contribute to a Green Fund aiming to support developing countries in their efforts to combat climate change.


Another interesting finding is that the reductions of emissions are much more evenly distributed across countries in C2 than in C1. Figure 4.5 shows that the average reductions in each country for 2030-2050 are all between 12 and 22 percent if they negotiate a carbon price, while they fluctuate widely between almost 55 percent and minus 33 percent if parties rely on nationally determined contributions. In sum, this leads to an overall lower actual average reduction over 2030-2050 in Committee 1 (about 5 percent) compared to a high reduction in Committee 2 (about 20 percent). Thus, despite the large contributions by four countries in C1, the carbon emission reductions are still below the reductions in C2.

Result 4: There are substantial and highly significant differences in the reductions of emissions for 2030-2050 between countries in C1, while these differences are much smaller and statistically not significant in C2.

Result 4 confirms Hypothesis 4. It is supported by nonparametric Kruskal-Wallis tests, which indicate significant differences of emission reductions between the countries in C1 ($H(9) = 103.15$, $p < .001$), but not in C2 ($H(9) = 1.53$, $p = .997$). We further analyze the differences for each country separately. In a first step, we compare the differences in 2030-2050 between C1 and C2 for each country. Supported by two-sided Wilcoxon signed-rank tests, we find that CO2 emission reduction goals are significantly different between C1 and C2 for all countries, except India and South Africa. In a second step, we run one-sided tests to further explore those eight countries showing significant differences. One-sided Wilcoxon signed-rank tests reveal that for

²³An often-heard argument by the delegates of Russia and Saudi Arabia was that the lower prices in 2030 would provide them with more time to adapt and hence also increase their willingness to accept the higher prices in the future.

Figure 4.5: Reduction of emissions of each country in Committee C1 and Committee C2 in 2030-2050 (averaged over all COP meetings).



the EU, Japan, Canada, and Australia, CO₂ emission reduction goals are significantly higher in C1 than in C2. On the contrary, the CO₂ emission reduction goals for China, Russia, Saudi Arabia, and the USA are significantly higher in C2 than in C1.

Result 4 is partially imposed by the negotiation design. In C2, countries have to agree to a uniform minimum price for carbon, so all countries that agree to the resolution automatically impose the same carbon price that yields the same reduction of emissions. During the negotiations in C2, some countries like the EU and Japan announced that they were willing to voluntarily impose a higher carbon price than the minimum carbon price agreed upon in the Resolution, but this was not accounted for in the actual emission reductions. If the negotiation protocol would have allowed for voluntary higher commitments in C2, this would have led to a less equal distribution of reductions and it might have further increased overall emissions in C2.²⁴

It is interesting to compare the experimental outcomes for the different countries to the real nationally determined contributions listed by these countries in the Annexes to the Paris agreement. However, countries have a lot discretion how to list and count their contributions, and not all countries submitted their NDC. Climate Action Tracker (CAT)²⁵ is an independent scientific analysis platform that tracks government climate action over time. Through its analyses and estimates, CAT helps to compare officially submitted NDCs. In the Online Appendix,

²⁴However, this could also reduce overall reductions, because countries that want a lower carbon price could argue that the minimum price should be kept low because those countries that want a higher price can have it on a voluntary basis.

²⁵See <https://climateactiontracker.org/about/>.

we compare the real NDCs of the Paris Agreement available at the time of our experiments (as measured by CAT) to the average behavior in C1 for 2030 and 2050. As can be seen in Table A.6, the large majority of the pledges in C1 are very close to the NDC counterparts. This confirms that the subjects took their roles in the simulated negotiations seriously.

4.6 Conclusion

MUN simulations are a novel experimental method that combines some of the advantages of laboratory and field experiments. They provide a formally structured framework that makes it possible to observe the outcomes of many simulated COP conferences on the same topic that were held under the exact same rules, with tight control over the preparation and the information provided to the participants. Thus, MUN conferences allow the experimenter to apply similarly high standards of replicability as in lab experiments, while at the same time enriching the experimental context and making the negotiations more realistic. This increases external validity. MUN simulations do not use monetary incentives but rely on the intrinsic motivation of the participants to engage with a complex topic and to represent ‘their’ nation as genuinely as possible. This may be closer to the motivation of real delegates than paying the subjects small amounts of money for a narrowly defined performance.²⁶ The lively and dynamic debating process allows researchers to account for the complexity of negotiation processes and to collect data on a variety of interesting variables, as for instance which arguments are used, which coalitions are formed (and abandoned), and which strategies are employed. These observations are insightful, but it is sometimes difficult to draw statistically validated conclusions from them. Of course, the new method also has its drawbacks. It requires a large organizational effort which limits the number of observations. Furthermore, the increase of external validity goes along with less control and thereby less internal validity of the experiment as compared to the lab.

Despite these drawbacks, MUN simulations can provide valuable empirical information in addition to lab experiments and case studies. Result 1 shows that negotiating a common commitment on a uniform carbon price is more successful in reducing carbon emissions in the long run than negotiating a non-binding common goal that has to be achieved by individual commitments, confirming the lab experiment of (Schmidt and Ockenfels (2021)). Thus, our experiment provides causal evidence that the negotiation setup affects the reductions of carbon emissions.

There are several other results that have important implications. First, Result 2 shows that the individual commitments (as in Paris style negotiations) do not live up to the non-binding common goal that the parties agreed upon. This is directly in line with projections of

²⁶See Voslinsky and Azar (2021) for a recent survey showing that experiments without financial incentives may be preferable to narrowly incentivized experiments.

the Paris agreement (e.g. Cléménçon (2016), Rogelj et al. (2016), Jiang et al. (2019), UNEP (2019) showing that the intended nationally determined pledged CO₂ emission reductions are not sufficient to reach the two-degree-goal. Result 3 shows that participation in the agreement is higher if a uniform carbon price is negotiated. This result is surprising as it is less costly for the nations to join a non-binding agreement, such as the Paris Agreement, as compared to a binding agreement, such as on a uniform global carbon price. However, a uniform carbon price limits free-riding because fossil fuel producers can leverage their participation for keeping the carbon price low.

Our final Result 4 shows that in C1 some countries (e.g. EU, Japan, Canada, and Australia) frequently reduced their emissions by more than the common goal, while other countries (e.g. Russia, Saudi Arabia, and the USA) contributed very little. The uniform carbon price in C2 forced all participating countries to contribute the same. Countries were not allowed to register additional voluntary climate action in the resolution. The discussions in C2 and the results of C1 showed, however, that some countries would have been prepared to do much more. This might have further increased the effectiveness of the carbon price negotiations.

Taken together our results lend some support for the proposal to focus international climate negotiations on a uniform carbon price. However, additional research is required to test the validity of this proposal. Furthermore, it is important to study how to best negotiate a Green Fund that is necessary to support developing countries in their efforts to combat and deal with climate change. These problems have been left aside in our study. Finally, it would be very interesting to study the dynamic features of different negotiation designs: Is it more difficult to achieve an upward spiral over time of individual commitments or of a uniform common commitment?

4.A Appendix I

The Online Appendix is provided on degruyter.com. It contains additional material and further analyses, including a comparison of nationally determined contributions from Climate Action Tracker to the average behavior in C1. Moreover, we provide the Study Guide, the QARMA of Committee 1 and Committee 2, and the Rules of Procedure as distributed to the delegates. Finally, the Online Appendix contains the Position Papers that were selected for an award.

Bibliography

- Angeletos, George-Marios, Fabrice Collard, and Harris Dellas (2018). “Quantifying Confidence”. In: *Econometrica* 86.5, pp. 1689–1726.
- (2020). “Business-Cycle Anatomy”. In: *American Economic Review* 110.10, pp. 3030–3070.
- Arias, Jonas E., Dario Caldara, and Juan F. Rubio-Ramírez (2019). “The Systematic Component of Monetary Policy in SVARs: An Agnostic Identification Procedure”. In: *Journal of Monetary Economics* 101, pp. 1–13.
- Arias, Jonas E., Juan F. Rubio-Ramírez, and Daniel F. Waggoner (2014). “Inference Based on SVARs Identified with Sign and Zero Restrictions: Theory and Applications Inference”. In: *Board of Governors of the Federal Reserve System International Finance Discussion Papers* 1100.
- (2018). “Inference Based on Structural Vector Autoregressions Identified With Sign and Zero Restrictions: Theory and Applications”. In: *Econometrica* 86.2, pp. 685–720.
- (2021). “Inference in Bayesian Proxy-SVARs”. In: *Journal of Econometrics* 225.1, pp. 88–106.
- Asal, Victor (2005). “Playing Games with International Relations”. In: *International Studies Perspectives* 6.3, pp. 359–373.
- Backus, David K. and Patrick J. Kehoe (1992). “International Evidence on the Historical Properties of Business Cycles”. In: *The American Economic Review* 82.4, pp. 864–888.
- Backus, David K., Patrick J. Kehoe, and Finn E. Kydland (1992). “International Real Business Cycles”. In: *Journal of Political Economy* 100.4, pp. 745–775.
- Bai, Yan, Jose-Victor Rios-Rull, and Kjetil Storesletten (2017). “Demand Shocks as Productivity Shocks”. In: *Unpublished*.
- Barrett, Scott (1994). “Self-Enforcing International Environmental Agreements”. In: *Oxford Economic Papers* 46, pp. 878–894.
- (2003). *Environment and Statecraft: The Strategy of Environmental Treaty-Making*. 1. publ. i. Oxford: Oxford University Press.

- Barrett, Scott (2005). “The theory of international environmental agreements”. In: *Handbook of environmental economics* vol. 3, pp. 1457–1516.
- Barrett, Scott and Astrid Dannenberg (2016). “An Experimental Investigation into ‘Pledge and Review’ in Climate Negotiations”. In: *Climatic Change* 138.1-2, pp. 339–351.
- Barsky, Robert B. and Eric R. Sims (2011). “News shocks and business cycles”. In: *Journal of Monetary Economics* 58.3, pp. 273–289.
- (2012). “Information, animal spirits, and the meaning of innovations in consumer confidence”. In: *American Economic Review* 102.4, pp. 1343–1377.
- Battaglini, Marco and Bård Harstad (2016). “Participation and Duration of Environmental Agreements”. In: *Journal of Political Economy* 124.1, pp. 160–204.
- (2020). “The Political Economy of Weak Treaties”. In: *Journal of Political Economy* 128.2, pp. 544–590.
- Bauer, Michael D. and Eric T. Swanson (2022). “A Reassessment of Monetary Policy Surprises and High-Frequency Identification”. In: *NBER Working Paper 29939*.
- Beaudry, Paul, Deokwoo Nam, and Jian Wang (2011). “Do Mood Swings Drive Business Cycles and is it Rational?” In: *NBER Working Paper Series* 2011.98.
- Beaudry, Paul and Franck Portier (2006). “Stock Prices, News, and Economic Fluctuations”. In: *The American Economic Review* 96.4, pp. 1293–1307.
- (2014). “Understanding noninflationary demand-driven business cycles”. In: *NBER Macroeconomics Annual*. Vol. 28. 1, pp. 69–130.
- Benati, Luca (2008). “Investigating Inflation Persistence across monetary regimes”. In: *The Quarterly Journal of Economics* 123: 3.August, pp. 1005–1060.
- (2021). “Leaning against House Prices: A Structural VAR Investigation”. In: *Journal of Monetary Economics* 118, pp. 399–412.
- (2022). “Exploring the Trade-Off Between Leaning Against Credit and Stabilizing Real Activity”. In: *Unpublished*.
- Benati, Luca, Robert E. Lucas, Juan Pablo Nicolini, and Warren Weber (2021). “International Evidence on Long-Run Money Demand”. In: *Journal of Monetary Economics* 117, pp. 43–63.
- Benhabib, Jess and Roger E.A. Farmer (1994). *Indeterminacy and increasing returns*.
- BIS (2014). “Bank for International Settlements, 84th Annual Report”. In: Basel: Bank for International Settlements.

- (2016). “Bank for International Settlements, 86th Annual Report”. In: Basel: Bank for International Settlements.
- Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry (2018). “Really Uncertain Business Cycles”. In: *Econometrica* 86.3, pp. 1031–1065.
- Boardman, Robert (1969). “The Theory and Practice of Educational Simulation”. In: *Educational Research* 11.3, pp. 179–184.
- Caldara, Dario and Edward Herbst (2019). “Monetary Policy, Real Activity, and Credit Spreads: Evidence from Bayesian Proxy SVARs”. In: *American Economic Journal: Macroeconomics* 11.1, pp. 157–192.
- Canova, Fabio and Matthias Paustian (2011). “Business cycle measurement with some theory”. In: *Journal of Monetary Economics* 58.4, pp. 345–361.
- Cléménçon, Raymond (2016). “The Two Sides of the Paris Climate Agreement”. In: *Journal of Environment and Development* 25.1, pp. 3–24.
- Comin, Diego and Mark Gertler (2006). “Medium-Term Business Cycles”. In: *The American Economic Review* 96.3, pp. 523–551.
- Comin, Diego, Javier Quintana, Tom Schmitz, and Antonella Trigari (2022). “Measuring TFP: The Role of Profits, Adjustment Costs, and Capacity Utilization”. In: *SSRN Electronic Journal* 72073.
- Dannenberg, Astrid and Carlo Gallier (2020). “The choice of institutions to solve cooperation problems: a survey of experimental research”. In: *Experimental Economics* 23.3, pp. 716–749.
- Dannenberg, Astrid, Andreas Lange, and Bodo Sturm (2014). “Participation and Commitment in Voluntary Coalitions to Provide Public Goods”. In: *Economica* 81.322, pp. 257–275.
- Dées, Stéphane and Jochen Güntner (2014). “The International Dimension of Confidence Shocks”. In: *European Central Bank Working Paper* N. 1669.16.
- Eichenbaum, Martin and Charles L Evans (1995). “Some Empirical Evidence on the Effects of Shocks to Monetary Policy on Exchange Rates”. In: *The Quarterly Journal of Economics* 110.4, pp. 975–1009.
- Elliott, B Y Grahwa, Thomas J Rothenberg, and James H Stock (1996). “Efficient Tests for an Autoregressive Unit Root”. In: *Econometrica* 64.4, pp. 813–836.
- Fève, Patrick and Alain Guay (2019). “Sentiments in SVARs”. In: *The Economic Journal* 129.618, pp. 877–896.

- Fisher, Jonas (2006). “The Dynamic Effects of Neutral and Investment - Specific Technology Shocks”. In: *Journal of Political Economy* 114.3, pp. 413–451.
- Fitzgerald, Terry J. and Lawrence J. Christiano (2003). “The Band Pass Filter”. In: *International Economic Review*, 44.2, pp. 435–465.
- Fry, Renée and Adrian Pagan (2011). “Sign restrictions in structural vector autoregressions: A critical review”. In: *Journal of Economic Literature* 49.4, pp. 938–960.
- Gertler, Mark and Peter Karadi (2015). “Monetary Policy Surprises, Credit Costs, and Economic Activity”. In: *American Economic Journal: Macroeconomics* 7.1, pp. 44–76.
- Gilchrist, Simon and Egon Zakrajšek (2012). “Credit spreads and business cycle fluctuations”. In: *American Economic Review* 102.4, pp. 1692–1720.
- Greenwood, By Jeremy, Zvi Hercowitz, and Gregory W Huffman (1988). “Investment, Capacity Utilization, and the Real Business Cycle”. In: *American Economic Review* 78.3, pp. 402–417.
- Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell (1997). “Long-Run Implications of Investment-Specific Technological Change”. In: *American Economic Review* 87.3, pp. 342–362.
- (2000). “The role of investment-specific technological change in the business cycle”. In: *European Economic Review* 44.1, pp. 91–115.
- Gürkaynak, Refet S, Brian Sack, and Eric T Swanson (2005). “Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements”. In: *International Journal of Central Banking* 1.1, pp. 55–93.
- Harrison, Glenn W and John A List (2004). “Field Experiments”. In: *Journal of Economic Literature* 42.4, pp. 1009–1055.
- Huo, Zhen, Andrei A. Levchenko, and Nitya Pandalai-Nayar (2020). “Utilization-Adjusted TFP Across Countries: Measurement and Implications for International Comovement”. In: *Working Paper*.
- IPCC (2014). *Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Ed. by Ottmar Edenhofer. Cambridge, UK and NY, USA: Cambridge University Press.
- Jiang, Xuemei, Glen P Peters, and Christopher Green (2019). “Global Rules Mask the Mitigation Challenge Facing Developing Countries”. In: *Earth’s Future* 7.4, pp. 428–432.
- Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti (2011). “Investment shocks and the relative price of investment”. In: *Review of Economic Dynamics* 14.1, pp. 102–121.

- Kauneckis, Derek L and Matthew R Auer (2013). "A Simulation of International Climate Regime Formation". In: *Simulation and Gaming* 44.2-3, pp. 302–327.
- Kocher, Martin G, Peter Martinsson, Emil Persson, and Xianghong Wang (2016). "Is there a hidden cost of imposing a minimum contribution level for public good contributions?" In: *Journal of Economic Psychology* 56, pp. 74–84.
- Koop, Gary (2003). *Bayesian Econometrics*. Hoboken, NJ: John Wiley & Sons.
- Kosfeld, Michael, Akira Okada, and Arno Riedl (2009). "Institution formation in public goods games". In: *American Economic Review* 99.4, pp. 1335–1355.
- Leeper, Eric M. and Tao Zha (2003). "Modest Policy Interventions". In: *Journal of Monetary Economics* 50.8, pp. 1673–1700.
- Lequiller, F. and D. Blades (2014). *Understanding National Accounts: Second Edition*.
- Lester, James P and Michael J Stoil (1979). "Evaluating a Role-Specific Simulation". In: *Simulation and Games* 10.2, pp. 167–181.
- Levchenko, Andrei A. and Nitya Pandalai-Nayar (2020). "TFP, News, and "Sentiments": The International Transmission of Business Cycles". In: *Journal of the European Economic Association* 18.1, pp. 302–341.
- List, John A. (2020). "Non est disputandum de generalizability? A glimpse into the external validity trial". In: *NBER Working Paper No. 27535* 27535.
- Lorenzoni, Guido (2009). "A theory of demand shocks". In: *American Economic Review* 99.5, pp. 2050–2084.
- Lucas, Robert E. (1977). "Understanding Business Cycles". In: *Carnegie-Rochester Conference Series on Public Policy* 5, pp. 7–29.
- MacKay, David J C, Peter Cramton, Axel Ockenfels, and Steven Stoft (2015). "Price Carbon - I Will if You Will". In: *Nature* 526.7573, pp. 315–316.
- Martinsson, Peter and Emil Persson (2019). "Public Goods and Minimum Provision Levels: Does the Institutional Formation Affect Cooperation?" In: *The Scandinavian Journal of Economics* 121.4, pp. 1473–1499.
- Matzner, Nils and Robert Herrenbrück (2017). "Simulating a Climate Engineering Crisis". In: *Simulation and Gaming* 48.2, pp. 268–290.
- Mertens, Karel and Morten O. Ravn (2013). "The dynamic effects of personal and corporate income tax changes in the United States". In: *American Economic Review* 103.4, pp. 1212–1247.

- Miyamoto, Wataru, Thuy Lan Nguyen, and Hyunseung Oh (2022). “In Search of Dominant Drivers of the Real Exchange Rate”. In: *Working Paper*, pp. 01–56.
- Muldoon, James P. (1995). “The Model United Nations Revisited”. In: *Simulation and Gaming* 26.1.
- Nam, Deokwoo and Jian Wang (2019). “Mood Swings and Business Cycles: Evidence from Sign Restrictions”. In: *Journal of Money, Credit and Banking* 51.6, pp. 1623–1649.
- Nordhaus, William (2015). “Climate Clubs: Overcoming Free-riding in International Climate Policy”. In: *American Economic Review* 105.4, pp. 1339–1370.
- (2019). “Climate Change: The Ultimate Challenge for Economics”. In: *American Economic Review* 109.6, pp. 1991–2014.
- Ohanian, Lee E. and Andrea Raffo (2012). “Aggregate hours worked in OECD countries: New measurement and implications for business cycles”. In: *Journal of Monetary Economics* 59.1, pp. 40–56.
- Penetrante, Ariel M. (2012). “Simulating Climate Change Negotiations: Lessons from Modeled Experience”. In: *Negotiation Journal* 28.3, pp. 279–314.
- Ramey, Valerie A. (2016). *Macroeconomic Shocks and Their Propagation*. 1st ed. Vol. 2. Elsevier B.V., pp. 71–162.
- Rogelj, Joeri, Michel den Elzen, Niklas Höhne, Taryn Fransen, Hanna Fekete, Harald Winkler, Roberto Schaeffer, Fu Sha, Keywan Riahi, and Malte Meinshausen (2016). “Paris agreement climate proposals need a boost to keep warming well below 2 °C”. In: *Nature* 534.
- Schelling, Thomas C. (1960). *The Strategy of Conflict*. Cambridge, Mass.: Harvard University Press.
- Schmidt, Klaus M. and Axel Ockenfels (2021). “Focusing climate negotiations on a uniform common commitment can promote cooperation”. In: *Proceedings of the National Academy of Sciences of the United States of America* 118.11.
- Schularick, Moritz, Lucas ter Steege, and Felix Ward (2021). “Leaning against the Wind and Crisis Risk”. In: *American Economic Review: Insights* 3.2, pp. 199–214.
- Schwardmann, Peter, Egon Tripodi, and Van der Weele, Joël J. (2022). “Self-Persuasion: Evidence from International Debating Competitions”. In: *American Economic Review* forthcoming.
- Sims, Christopher A. (1980). “Macroeconomics and Reality”. In: *Econometrica* 48.1, pp. 1–48.
- Smets, Frank and Rafael Wouters (2007). “Shocks and frictions in US business cycles: A Bayesian DSGE approach”. In: *American Economic Review* 97.3, pp. 586–606.

- Stiglitz, Joseph E., Nicholas Stern, Maosheng Duan, Ottmar Edenhofer, Gaël Giraud, Geoffrey M. Heal, Emilio Lèbre La Rovere, Adele Morris, Elisabeth Moyer, Mari Pangestu, Priyadarshi R. Shukla, Youba Sokona, and Harald Winkler (2017). *Report of the High-Level Commission on Carbon Prices*. New York.
- Stock, James H. and Mark W. Watson (2012). “Disentangling the Channels of the 2007-09 Recession”. In: *Brookings Papers on Economic Activity* 1, pp. 81–135.
- (2016). *Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics*. 1st ed. Vol. 2. Elsevier B.V., pp. 415–525.
- Svensson, Lars E.O. (2017). “Cost-Benefit Analysis of Leaning against the Wind”. In: *Journal of Monetary Economics* 90, pp. 193–213.
- Uhlig, Harald (1998). “The Robustness of Identified VAR Conclusions about Money”. In: *International Finance Discussion Paper* 1998.610, pp. 1–50.
- (2003). “What moves real GNP?” In: *Unpublished manuscript*.
- (2004). “Do technology shocks lead to a fall in total hours worked?” In: *Journal of the European Economic Association* 2.2-3, pp. 361–371.
- (2005). “What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure”. In: *Journal of Monetary Economics* 52.2, pp. 381–419.
- UNEP (2019). “Emissions Gap Report 2019”. In: *United Nations Environment Programme*.
- UNFCCC Secretariat (2015). “Synthesis Report on the Aggregate Effect of the Intended Nationally Determined Contributions”. In: *UN Framework Convention on Climate Change*. FCCC/CP/2015/7. Ed. by UN Framework Convention on Climate Change.
- Voslinsky, Alisa and Ofer H Azar (2021). “Incentives in experimental economics”. In: *Journal of Behavioral and Experimental Economics* 93, p. 101706.
- Waggoner, Daniel F. and Tao Zha (2003). “A Gibbs Sampler for Structural Vector Autoregressions”. In: *Journal of Economic Dynamics and Control* 28.2, pp. 349–366.
- Weitzman, Martin L. (2014). “Can Negotiating a Uniform Carbon Price Help to Internalize the Global Warming Externality?” In: *Journal of the Association of Environmental and Resource Economists* 1.1/2, pp. 29–49.
- (2015). “Internalizing the Climate Externality: Can a Uniform Price Commitment Help?” In: *Economics of Energy and Environmental Policy* 4.2, pp. 37–50.
- (2017a). “On a world climate assembly and the social cost of carbon”. In: *Economica* 84.336, pp. 559–586.

Weitzman, Martin L. (2017b). “Voting on Prices vs . Voting on Quantities in a World Climate Assembly”. In: *Research in Economics* 71.2, pp. 199–211.

Statement of Authorship

I declare herewith that I wrote this thesis on my own, without the help of others. Wherever I have used permitted sources of information, I have made this explicitly clear within my text and I have listed the referenced sources. I understand that if I do not follow these rules that the Senate of the University of Bern is authorized to revoke the title awarded on the basis of this thesis according to Article 36, paragraph 1, litera r of the University Act of September 5th, 1996.

Selbständigkeitserklärung

Ich erkläre hiermit, dass ich diese Arbeit selbständig verfasst und keine anderen als die angegebenen Quellen benutzt habe. Alle Koautorenschaften sowie alle Stellen, die wörtlich oder sinngemäss aus Quellen entnommen wurden, habe ich als solche gekennzeichnet. Mir ist bekannt, dass andernfalls der Senat gemäss Artikel 36 Absatz 1 Buchstabe o des Gesetzes vom 5. September 1996 über die Universität zum Entzug des aufgrund dieser Arbeit verliehenen Titels berechtigt ist.

Bern, 31. Januar 2023

Lucas Christopher Andreas Kyriacou