

Disentangling Macroeconomic Policies

José Miguel Acosta

Submitted in partial fulfillment of the  
requirements for the degree of  
Doctor of Philosophy  
under the Executive Committee  
of the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2022

© 2022

José Miguel Acosta

All Rights Reserved

## **Abstract**

### Disentangling Macroeconomic Policies

José Miguel Acosta

The field of macroeconomics has increasingly turned its attention towards understanding the state dependent effects of macroeconomic policies, the idea being that different macroeconomic conditions—e.g., the state of the business cycle, or the distribution of income over the population—can cause a single economic policy shock to propagate differently through the economy. I turn this thinking around in this dissertation, and instead ask whether the policies that we study are, in practice, “single economic policy shocks” or are, instead, aggregates of multiple policies with different effects. In chapter 1, I decompose monetary policy into an interest rate component, and a component that captures macroeconomic information provision. In chapter 2, I discuss the consequences of tariff policy in light of the view that tariffs in the United States are extremely heterogeneous and, on average, regressive in nature. In chapter 3, I return to monetary policy, asking whether more-transparent communications from the Federal Reserve have allowed for a more effective transmission of monetary policy. To summarize my findings, in all cases I find that the economic consequences estimated using disaggregated policy measures differ substantially from the consequences estimated using aggregated measures.

## Table of Contents

Acknowledgments . . . . .	xi
Dedication . . . . .	xii
Preface . . . . .	1
Chapter 1: The Perceived Causes of Monetary Policy Surprises . . . . .	4
1.1 Introduction . . . . .	4
1.2 Data: Construction and Validation . . . . .	11
1.2.1 High-Frequency Text-Based Proxies of Macroeconomic Expectations . . . . .	12
1.2.2 Validation . . . . .	17
1.2.3 Predictability . . . . .	22
1.2.4 Numerical Data . . . . .	24
1.3 Shock Identification . . . . .	25
1.3.1 The Identification Challenge in an Illustrative Model . . . . .	26
1.3.2 Regimes and Intuition for the Identification Assumptions . . . . .	30
1.3.3 Formal Identification Assumptions . . . . .	33
1.3.4 Estimation . . . . .	38
1.4 Effects of Monetary Policy and Information . . . . .	43
1.4.1 Baseline Results . . . . .	43

1.4.2	Comparison with Previous Estimates . . . . .	46
1.4.3	Comparison with Previous Estimates: Magnitudes . . . . .	50
1.5	Expanding the Picture: Short-term Rates and Supply Factors . . . . .	52
1.5.1	Short- vs. Long-term Interest Rates . . . . .	53
1.5.2	Supply vs. Demand Factors . . . . .	57
1.6	Conclusion . . . . .	60
Chapter 2: The Regressive Nature of the U.S. Tariff Code: Origins and Implications . . . . .		62
2.1	Introduction . . . . .	62
2.2	Documenting the Pattern . . . . .	66
2.2.1	Defining a “Good” . . . . .	67
2.2.2	Coverage of the Tariff Schedule . . . . .	70
2.2.3	Regressivity Across the Tariff Code . . . . .	71
2.2.4	Robustness: Alternative Definitions and Global Presence . . . . .	74
2.2.5	Regressivity is Not a New Phenomenon . . . . .	75
2.3	Origins of Regressivity: A Case Study Approach . . . . .	77
2.3.1	Concessions and the Case of Fishing Reels . . . . .	78
2.3.2	Protection and the Cases of Forks and Bicycles . . . . .	82
2.3.3	Summary: Regressivity Came from Bygone Policy Priorities . . . . .	84
2.4	New Data on Legislated Tariff Rates . . . . .	85
2.5	Origins of Regressivity: Systematic Analysis . . . . .	88
2.5.1	Hysteresis in A Long Panel of Consumer Goods . . . . .	90
2.5.2	The Introduction of Regressivity: 1930–1946 . . . . .	92

2.5.3	Within-Good Tariff Variation Post 1946 . . . . .	96
2.6	Economic Significance . . . . .	97
2.6.1	Aggregate Implications . . . . .	98
2.6.2	Sufficient Statistic Calculation . . . . .	99
2.6.3	Data . . . . .	100
2.6.4	Estimation . . . . .	102
2.7	Conclusion . . . . .	103
Chapter 3:	A New Measure of Central Bank Transparency and Implications for the Effec- tiveness of Monetary Policy . . . . .	105
3.1	Introduction . . . . .	105
3.2	Measuring Procedural Transparency . . . . .	107
3.2.1	Previous Measurements . . . . .	108
3.2.2	How to Measure Transparency . . . . .	108
3.2.3	Language Model . . . . .	110
3.2.4	Transparency Index Definition . . . . .	115
3.3	Notable Features of the Minutes, Transcripts, and Procedural Transparency . . . . .	116
3.3.1	Correlations with External Variables . . . . .	117
3.3.2	Anecdotal Understanding of the Measure . . . . .	120
3.3.3	Observability of the Measure . . . . .	121
3.4	Effectiveness of Monetary Policy . . . . .	123
3.4.1	Defining Effectiveness . . . . .	123
3.4.2	Empirical Strategy . . . . .	124
3.4.3	Results and Robustness . . . . .	125

3.4.4	Omitted Variables Bias . . . . .	128
3.5	Interpretation: Transparency about What? . . . . .	129
3.5.1	Economic Outlook and Policy Strategy Transparency . . . . .	130
3.5.2	Topic-Specific Transparency . . . . .	131
3.6	Conclusion . . . . .	135
References . . . . .		136
Appendix A: Numerical Example: Textbook New Keynesian Model . . . . .		146
Appendix B: Formal Signaling Structures . . . . .		149
B.1	Static Signaling Model . . . . .	149
B.2	Dynamic Extension . . . . .	150
Appendix C: Solving the Noisy Information Model . . . . .		154
Appendix D: Why is $\hat{y}_t$ an invalid instrument (and bad control)? . . . . .		157
Appendix E: Text Analysis Appendix . . . . .		160
E.1	Word Lists . . . . .	160
E.2	Word List Construction . . . . .	161
E.3	Machine Learning Approach . . . . .	161
Appendix F: Alternative “Second Variables” . . . . .		165
Appendix G: Estimation and Bootstrap Procedure . . . . .		168
Appendix H: Local Projections: Robustness . . . . .		174

Appendix I: Predictability of Interest-Rate Surprises: A Comment on Bauer and Swanson (2020)	178
Appendix J: Assessing the Empirical Assumptions	180
J.1 Uncorrelated Forecast Revisions	180
J.2 Constant Mapping between Shocks and Observables	182
J.2.1 Statistical Validity	182
J.2.2 Empirical Validity: Melosi (2017)	182
J.2.3 Dispersed vs. Common Knowledge	189
Appendix K: Weighted Statistics	191
Appendix L: Regressivity in the European Tariff Schedule	192
Appendix M: Details on Digitized Tariff Data	193
Appendix N: Transcripts, Minutes, ROPAs, MOAs, and MODs: The History of FOMC Communications	196
Appendix O: Definition and Relevance of Procedural Transparency	199
Appendix P: Number of Topics: Cross-Validation and Robustness	202
Appendix Q: Robustness for Moving Averages	203
Appendix R: Tabular Representation of Figures 3.4 and 3.5	204
Appendix S: Robustness to Text-Analysis Features	206

## List of Figures

1.1	Newspaper-based Output Directionality and Macroeconomic Expectations . . . . .	15
1.2	High-frequency Forecast Revisions of Output and Interest Rates . . . . .	16
1.3	Blue Chip GDP Forecast Revisions and the High-frequency Proxy . . . . .	17
1.4	Shapley Decomposition of Forecast Revision $R^2$ . . . . .	21
1.5	The Identification Assumptions in Pictures . . . . .	31
1.6	Time Series of Estimated Monetary and Information Shocks . . . . .	40
1.7	Variance Decomposition of HF Forecast Revisions . . . . .	41
1.8	Baseline Macroeconomic Effects . . . . .	42
1.9	One-Day Changes in Yields . . . . .	44
1.10	Effects of Monetary Shocks: Using Longer-term Interest Rates . . . . .	46
1.11	Response of Industrial Production to Traditional Contractionary MP Shocks . . . . .	48
1.12	Response of Industrial Production to “Information Free” Contractionary MP Shocks . . . . .	50
1.13	LP-IV Estimates . . . . .	51
1.14	Federal Funds Rate Surprise and the Target Shock . . . . .	55
1.15	High-frequency Text-based Inflation Expectations . . . . .	57
2.1	Heterogeneity and Overall Presence of the Pattern . . . . .	72
2.2	Tariff Differentials between Low- and High-Value Goods . . . . .	73

2.3	Legislated Tariffs, 1989-2017 . . . . .	77
2.4	History of Tariff Rates on Fishing Reels . . . . .	80
2.5	History of Tariff Rates on Stainless Steel Forks . . . . .	82
2.6	Examples of Digitized Tariff Schedules: Manganese Compounds . . . . .	87
2.7	Tariff Rates Over Time for 15 Goods . . . . .	89
2.8	Average AVE Tariff Relative to 1930 for 15 Goods . . . . .	90
2.9	Source Countries for 15 Goods . . . . .	91
2.10	Tariffs on Goods Defined by Unit Value: Smoot-Hawley Tariff . . . . .	95
2.11	Tariff Changes in Valued Goods during 1930–1946 . . . . .	96
2.12	Within-Good Tariff Variation Share Relative to 1930 . . . . .	98
2.13	Revenue Loss from Eliminating Regressive Pattern . . . . .	99
2.14	Welfare Gains under Alternative Counterfactuals . . . . .	103
3.1	Procedural Transparency . . . . .	116
3.2	Entropy of the Minutes (Left) and Transcripts (Right) . . . . .	117
3.3	Monetary Policy Transparency in the News . . . . .	122
3.4	Monetary Policy Effectiveness: Regression Results . . . . .	126
3.5	Monetary Policy Effectiveness: Robustness of Regression Results . . . . .	127
3.6	Transparency by Part of Transcript . . . . .	131
3.7	Monetary Policy Effectiveness: Regression Results using MPS Similarity . . . . .	132
3.8	Monetary Policy Effectiveness: Regression Results using ECSIT Similarity . . . . .	133
3.9	Monetary Policy Effectiveness: Regression Results using Topics-Specific Similarities	134
F.1	Correlation Matrix of Variables Responding to Fed Announcements . . . . .	165

J.1	Heteroskedasticity in the Model . . . . .	185
J.2	Composition of the Heteroskedasticity Estimator . . . . .	187
M.1	Legislated Tariffs: Summary Statistics . . . . .	195
O.1	Google Trends: Fed Communications . . . . .	200
P.1	Cross-Validation Results . . . . .	202
Q.1	Estimates of $\gamma$ for Different Moving Averages . . . . .	203
S.1	Robustness of Transparency Measure to $K$ and Sample Period . . . . .	206
S.2	Robustness of Regression Results to $K$ and Sample Period . . . . .	207

## List of Tables

1.1	Low- and High-frequency Forecast Revisions: Robustness . . . . .	19
1.2	Predictability of High-frequency Output ( $\hat{y}_t$ ) and Interest-rate ( $\hat{i}_t$ ) Forecast Revisions	23
1.3	Heteroskedasticity-Based Estimates . . . . .	39
1.4	Shocks and Macroeconomic Expectations . . . . .	47
1.5	Identifying a Target Shock . . . . .	54
1.6	Comparison of Newspaper-Based Inflation Expectations with Breakeven Inflation .	59
2.1	2019 HTS: All Varieties with HTS Code 8215 . . . . .	68
2.2	Good-Level Summary Statistics . . . . .	70
2.3	Establishing the Pattern: Robustness . . . . .	76
2.4	GATT Participant and Concentration of High-Value Exports (1974) . . . . .	81
2.5	Income and Import Habits . . . . .	101
3.1	Topic Descriptions . . . . .	113
3.2	Regression Coefficients for Communications Variables . . . . .	118
3.3	Jointly Testing Correlates of Communications Measures . . . . .	119
E.1	Extended GDP Word List . . . . .	162
J.1	Allowing for Correlation in Perceived Shocks . . . . .	181

M.1 Data Sources . . . . . 193

N.1 FOMC Publications: Banking Act of 1935 to Present (2015). . . . . 196

O.1 The types of transparency relevant to central banks (Geraats, 2002, pp. F540) . . . 199

## Acknowledgements

This dissertation has benefited enormously from conversations with a host of advisors, collaborators, friends, family, and colleagues. I acknowledge contributions to specific chapters within the relevant chapter.

I am indebted to the members of my dissertation committee—Hassan Afrouzi, Jennifer La’O, Emi Nakamura, Jón Steinsson, and Michael Woodford—for pushing me to think carefully and rigorously about my research, while at the same time pushing me to communicate that research clearly. Their preparation started well before the research phase of graduate school: all but Hassan were my formal classroom instructors, though I think Hassan should qualify given his generosity in teaching me the ins and outs of models with imperfect information.

As hard as it is for me to believe, most of my life was *not* spent in graduate school. Before then, my research interests were fostered by Ellen Meade, John Taylor, and the 2014–2016 members of the Monetary Studies Unit at the Federal Reserve Board. None of this would have gotten off the ground without them.

I am extremely fortunate to have grown up with the unconditional love and support of my parents and sister: José Acosta, Mary Grebenc, and Lourdes Acosta. That support continued throughout graduate school, and made it a lot easier to get through the last six years. For all of that, and for having to partially experience graduate school, I will be forever grateful.

Finally, I am tremendously lucky to have found a life partner in Lydia Cox, whose role in making this dissertation possible is almost impossible to overstate. Lyd, I’m not sure what we’ve gotten ourselves into, but it will probably work.

## **Dedication**

To my grandparents: Deida Acosta Velez, Ramon Delfín Acosta, Dolores Grebenc, Joseph Grebenc, a.k.a Abuelita, Coco, Grandma, and Grandpa.

## **Introduction**

The field of macroeconomics has increasingly turned its attention towards understanding the state dependent effects of macroeconomic policies, the idea being that different macroeconomic conditions—e.g., the state of the business cycle, or the distribution of income over the population—can cause a single economic policy shock to propagate differently through the economy. I turn this thinking around in this dissertation, and instead ask whether the policies that we study are, in practice, “single economic policy shocks” or are, instead, aggregates of multiple policies with different effects. In chapter 1, I decompose monetary policy into an interest rate component, and a component that captures macroeconomic information provision. In chapter 2, I discuss the consequences of tariff policy in light of the view that tariffs in the United States are extremely heterogeneous and, on average, regressive in nature. In chapter 3, I return to monetary policy, asking whether more-transparent communications from the Federal Reserve have allowed for a more effective transmission of monetary policy. To summarize my findings, in all cases studied herein, I find that the economic consequences estimated using disaggregated policy measures differ substantially from the consequences estimated using aggregated measures.

In chapter 1, I address one of the central questions in macroeconomics: what are the macroeconomic effects of monetary policy? A well-known challenge to addressing this question is that, in practice, “monetary policy” is a multidimensional signal about both the path of interest rates, and the state of the economy more broadly. This makes it difficult to disentangle, or identify, the effects of monetary policy using only interest rates—the standard approach in the

literature. The first part of chapter 1 shows formally that, despite this practical complication, the identification challenge is no different than the standard problem of simultaneous determination that has been at the heart of empirical macroeconomics for decades. The framework contained therein allows for a recasting of the various approaches that have been used to address the identification challenge in these terms, so that the strengths and weaknesses of each approach can be evaluated according to established reasoning. Guided by the theoretical framework, in the second part of chapter 1, I introduce new data and identifying assumptions in order to separately identify the effects of communications about interest rates, and of communications about the state of the economy. These two aspects of policy have opposite effects on macroeconomic variables, for the same effect on interest rates. Studies that therefore only rely on interest rates (or take an approach that does not sufficiently reduce measurement error) tend to find minimal effects of changes in interest rates. Instead, I find much larger effects once I take into the account the actual aspect of policy that causes the change in rates.

Chapter 2 also has, at its center, a central macroeconomic policy question: what are the consequences of reducing barriers to trade. That chapter, written jointly with Lydia Cox, highlights that while models of international trade typically construe of tariffs as a homogeneous object, in practice, the U.S. tariff schedule is tremendously complex—today’s schedule comprising 4,394 pages of tariffs on 19,347 varieties of goods. Adding to this complexity, tariffs are not set uniformly across goods or industries, but are set at the *variety*-level. The monolithic nature of the tariff code nearly guarantees that individual tariff rates are not re-optimized each period, but are hysteretic: in many cases, tariff rates reflect tariff policies of bygone eras that were set with era-specific objectives in mind. With that in mind, in chapter 2, we document a fact in today’s tariff schedules that is emblematic of the hysteretic nature of the U.S. tariff code: tariffs are regressive in nature, with tariffs systematically higher on low-value versions of goods relative to their high-value counterparts. For example, the tariff on a \$400 handbag made of reptile leather is 5.3 percent, while the tariff on an \$8 plastic-sided handbag is 16 percent. To this point, this pattern has only been noted anecdotally. We show that this pattern is present across the entire

tariff schedule, and is especially pervasive among consumer goods. Using newly-digitized tariff schedules going back to the 1930 Smoot-Hawley Act, we show that regressivity originated in the 1930s and 40s and has persisted over time. This pattern is one that arises from “within good” tariff variation—i.e., tariffs on lower and higher end varieties of goods are set differently. Our data allow us to show that, across the entire tariff schedule, most within good variation in tariffs originated in the few decades following Smoot-Hawley, and persisted to today. While regressivity may have been a cost of policy objectives of the 1930s and 40s, we show that, today, this pattern substantially alters the implied distributional gains from trade.

In chapter 3, I return to an application in monetary economics, and ask whether more-transparent communications allow monetary policymakers to more effectively influence longer-term interest rates—an important policy tool, especially when interest rates are constrained. While transparency has been posited as a channel through which monetary policy is made more effective in this way, empirical studies of this question have lacked access to a time-varying high-frequency measure of transparency. To that end, in chapter 3, I present a new measure of the transparency of Federal Reserve deliberations, derived from the documents that the Fed uses to record and summarize each of its meetings. The measure—the similarity of the minutes and transcripts of each Federal Open Market Committee (FOMC) meeting—is largely, though not entirely, shaped by FOMC leadership. Monetary policy shocks have about a 40 percent larger effect on nominal and real interest rates when delivered transparently, suggesting an important role for transparency in determining the efficacy of monetary policy. These effects are primarily driven by transparency about monetary policy strategies conditional on the state of the economy.

# Chapter 1: The Perceived Causes of Monetary Policy Surprises

## 1.1 Introduction

The role of the Federal Reserve (Fed) has moved far beyond setting the level of the overnight interest rate. In the eyes of financial market participants, the Fed Chair is often seen as a fortune teller who communicates predictions regarding not only the path of future interest rates, but also forecasts of macroeconomic outcomes.<sup>1</sup> An understanding of the Fed's role in the economy, then, crucially hinges on an understanding of the macroeconomic effects of both aspects of Fed communications. Two challenges arise when estimating these effects. First, communications regarding the path of interest rates and communications regarding the economic outlook are both highly endogenous with respect to economic fundamentals. Second, the Fed typically engages in both types of communications simultaneously, which complicates attempts to separately identify their effects.

In this paper I estimate the macroeconomic consequences of Fed communications—both the effects of monetary policy (explicitly communicating about future interest rates) and information provision policy (the effects of providing information about macroeconomic fundamentals). I contribute to a longstanding literature that studies the effects of monetary policy using market reactions to Fed policy announcements, and provide new estimates that overcome an important conceptual

---

<sup>1</sup>In addition to the members of my dissertation committee, this chapter has also benefited from discussions with Andrea Ajello, Paul Bouscasse, Lydia Cox, Emilien Gouin-Bonenfant, Juan Herreño, Isaiah Hull (discussant), Ellen Meade, José Luis Montiel-Olea, Christian Moser, Joseph Saia, Jesse Schreger, Yeji Sung, Pierre Yared, and seminar participants at Berkeley, Columbia, Michigan, the Federal Reserve Board, the Federal Reserve Banks of Boston and Dallas, the 14<sup>th</sup> International Conference on Computational and Financial Econometrics, and the 2021 EEA Annual Congress. Part of this work was completed during a dissertation fellowship at the Federal Reserve Bank of San Francisco, although my discussions with the economists there were instrumental, the views expressed herein are mine and not those of the Fed. I owe a special thanks to Yasmin Saira at the Business and Economics Library at Columbia for help accessing data, especially when physical access to Columbia's libraries was limited due to the pandemic. Hannah Bingley and Blaine Helleloid provided outstanding research assistance. The work was supported by financial support from an Alfred P. Sloan Foundation Pre-Doctoral Fellowship in Behavioral Macroeconomics awarded through the NBER.

issue in the identification of these effects: Market-based measures of interest-rate expectations can respond to both types of policy, thereby identifying neither. I also take the view that information provision is an important component of the Fed’s communication policy—rather than a statistical nuisance that challenges the identification of exogenous variation in monetary policy, as the literature has come to perceive it.

To estimate the effects of both aspects of Fed communications on macroeconomic outcomes, I identify two new series of shocks: perceived monetary policy shocks (shocks to the policy rule) and “information shocks” (shocks to beliefs about economic fundamentals).<sup>2</sup> I posit that an econometrician needs access to (at least) two measures of market reactions that respond differently to the two types of shocks in order to identify them separately. I show that interest-rate and GDP forecast revisions emerge as natural candidates to accomplish the task, based on the implications of standard New Keynesian theory. Intuitively, because interest-rate and GDP forecast revisions react differently in response to monetary policy and information shocks, observing their joint reactions to monetary policy announcements can provide useful information for identifying the prevalence of each shock.<sup>3</sup>

The foundation of my shocks consists of two high-frequency measures of macroeconomic forecast revisions: interest-rate surprises and a new, text-based measure of GDP forecast revisions. The former is standard in the empirical monetary literature. The latter I construct using newspaper articles written about each Fed policy meeting. Specifically, I compute a GDP “directionality index” for articles written in a one-day window around each announcement. The index is based on the difference between increasing and decreasing mentions of GDP. I anticipate that this high-frequency series of GDP expectations will be useful in other contexts. My proxy for GDP forecast revisions is the unpredictable component of post-meeting directionality vis-à-vis pre-meeting directionality.

---

<sup>2</sup>I introduce the word “perceived” here because my shocks are based on data on expectations (as are estimates of monetary shocks in the literature that use interest rate futures data). The illustrative model in section 1.3.1 makes this observation precise.

<sup>3</sup>To give a familiar example, a contractionary monetary policy shock raises interest rates and lowers real GDP in standard New Keynesian models. Instead, a positive aggregate demand shock (which is what my information shocks resemble) causes both variables to increase.

I assume that macroeconomic expectations measured shortly before each policy announcement reflect the Fed’s communication expected by markets, given all macroeconomic events that have occurred up to the announcement. The difference between pre- and post-meeting expectations, then, should only arise from exogenous policy or information asymmetries. As such, the high-frequency construction of my variables addresses the general endogeneity of interest-rate and GDP forecasts with respect to observable economic fundamentals.

Because the Fed communicates simultaneously about interest rates and the macroeconomic outlook, I estimate a simultaneous-equations model of how markets update their interest-rate and GDP forecasts in response to Fed announcements, in order to recover the structural shocks from these forecast revisions. In the model, market participants are Bayesian forecasters whose model of the economy is a linear relationship between macroeconomic shocks (here, monetary policy and information shocks) and macroeconomic variables (here, interest rates and GDP). Despite this simple formulation—a system of two equations determined by two exogenous shocks—this is the forecasting model implied by the dynamic stochastic general equilibrium models that permeate macroeconomic analysis.

To identify the model, I use a discrete change in the Fed’s communication policy: the introduction of interest-rate forward guidance in 2003. As highlighted by Lunsford (2020), prior to 2003 the Fed’s post-meeting policy statements primarily described the economic outlook. In August 2003 the Fed began the practice of interest-rate forward guidance when it promised to keep interest rates low “for a considerable period.” To see how this can help identify the model, consider the implications of this policy change. Before 2003, learning about the Fed’s economic outlook was straightforward; in contrast, inference about the path of interest rates was possible only indirectly through the Fed’s discussion of the economy. This observation allows me to make my formal identification assumption: Fed announcements induced market participants to update their expectations about the future path of interest rates, relative to their expectations about economic fundamentals, more completely after 2003. My data support this assumption: Before 2003, interest-rate and GDP forecast revisions were positively correlated. Intuitively, this regime primarily provides identifica-

tion of the information shock. After 2003, the series are essentially uncorrelated. Therefore, this latter regime provides identification of the monetary shock by providing data whose variation is driven by a factor orthogonal to the information shock. This is “identification by heteroskedasticity,” proposed by Rigobon (2003).<sup>4</sup> Formally, while the model is not identified within either regime separately, the added assumption that the shocks have the same effects on observables across the two regimes imposes enough parametric restrictions to jointly solve the model’s implied moment conditions.<sup>5</sup> Importantly, while I estimate the model using a 3 year window around the 2003 policy change, I use the identified model to construct estimates of the two shocks over my entire 1999–2019 sample period.

I find that the monetary policy shock I identify has effects on macroeconomic outcomes and expectations that are consistent with New Keynesian macroeconomic models, a conclusion I reach without imposing such consistency a priori. A monetary shock that raises longer-term interest-rate expectations on impact leads to declines in industrial production and inflation. The responses of both variables are fairly delayed, with peak responses estimated between 2 and 3 years after the shock. Nominal and real interest rates increase on impact, while GDP and inflation expectations decrease. Notably, I use simple empirical specifications to identify these effects. When I substitute my measures for other estimates of contractionary monetary shocks in these specifications, the estimated responses of output are generally not in line with predictions from New Keynesian models. In terms of magnitudes, my results suggest that the effects of monetary policy are “big,” in the language of Coibion (2012), and similar to those estimated by Romer and Romer (2004): a 25 basis point increase in the policy rate leads to a roughly one percent decrease in industrial production.

Moving to the effects of information provision, I find that information shocks have effects on macroeconomic outcomes and expectations that are similar in (absolute) magnitude to those

---

<sup>4</sup>Other papers (Wright, 2012; Arai, 2017; Nakamura and Steinsson, 2018) have used the heteroskedasticity-based identification assumptions in seeking to estimate the effects of monetary policy. The approach is typically seen as a method of purging “background noise” (or latent factors, in the case of Gürkaynak et al. (2020)) from OLS regressions. I use the approach completely differently in that I am interested in estimating two shocks, not simply purging one shock of interest of a nuisance component. Lewis (2019) identifies similarly-named shocks using a heteroskedasticity-based approach that provides only statistical identification.

<sup>5</sup>This language comes from Stock and Watson (2017).

of monetary policy shocks. This remains true when controlling for recent macroeconomic news, which suggests first that the Fed plays an important role in its characterization of macroeconomic shocks.<sup>6</sup> In addition, my information shock creates a positive comovement of output and inflation (and expectations thereof, in low and high frequency), which suggests that the information primarily concerns demand-type factors.

The notion that monetary policy announcements can convey macroeconomic information, and thus contaminate estimates of exogenous monetary shocks, was put forth by Romer and Romer (2000). Campbell et al. (2012) and Nakamura and Steinsson (2018) highlight the fact that the presence of “information effects” can contaminate traditional high-frequency estimates of monetary policy shocks—my results confirm that the contamination is substantial: about 80% of a widely used high-frequency shock series is made up of information effects. Therefore, the empirical challenge became: How to disentangle information provision from exogenous monetary policy? My model shows that this challenge is the familiar problem of simultaneous determination encountered in supply and demand systems or structural vector autoregressions. Seen in this context, early work in this area imposed zero restrictions to identify monetary shocks, which are not warranted under the presence of information effects.<sup>7</sup> My work relaxes these restrictions. Nakamura and Steinsson (2018) estimate a structural model to overcome the identification challenge. In contrast to my work, their approach depends on all of the assumptions underlying the particular macroeconomic model.

Other papers cognizant of information effects have sought model-free approaches to identify high-frequency monetary policy shocks. Given that information effects are posited to stem from the Fed’s private information, Miranda-Agrippino and Ricco (2021) and Handlan (2020) propose orthogonalizing high-frequency interest-rate surprises to the Fed’s private information as captured

---

<sup>6</sup>This also addresses the critique of Bauer and Swanson (2020), whereby the positive correlation between interest rate surprises and macroeconomic expectations can result from both series’ reaction to the same economic news.

<sup>7</sup>This is the case in Kuttner (2001) (who introduced surprises in the current-meeting interest rate) and Gürkaynak et al. (2005) (who introduced the notion of a shock to the path of interest rates). By not entertaining the possibility that changes in interest-rate futures could be driven by shocks other than exogenous monetary policy, the authors imposed the restriction that other shocks had zero effect on high-frequency interest-rate changes.

by the Fed staff’s presentation materials (“Greenbook forecasts”).<sup>8</sup> This approach suffers from two conceptual shortcomings that I sidestep. First, it assumes that the staff’s economic assessment spans that of the Fed’s policymaking committee, which is ultimately tasked with policy communication. Romer and Romer (2008) show that these assessments generally do not align. Second, this approach requires choosing a set of variables that completely span the Fed’s private information *when the announcement is made*. I avoid having to posit the variables over which the Fed has private information, which is difficult to know, given the vast number of indicators that inform Fed policy decisions.<sup>9</sup> In addition, by relying on staff-created reports, neither paper can control for events that occur shortly before policy announcements, and thus both papers potentially fail to control for endogenous macroeconomic events to which the Fed might respond. My high-frequency measures avoid this concern.

Cieslak and Schrimpf (2019) and Jarociński and Karadi (2020) assume theoretically motivated sign restrictions regarding the relationship between monetary shocks, information shocks, stock returns, and interest-rate surprises. These sign restrictions allow the authors to discuss the relative importance of—and identify, in the case of Jarociński and Karadi (2020)—monetary and information shocks. In high frequency, stock returns and interest-rate surprises are consistently negatively correlated, in contrast to my estimates of output expectations. This suggests a limited role for stock prices in differentiating between monetary and information shocks. In fact, it suggests a limited role for information effects.<sup>10</sup> Examining a variable whose response to each shock differs sub-

---

<sup>8</sup>Campbell et al. (2012) also attempt to control for information effects by controlling for professional macroeconomic forecasts. As Woodford’s comment to that article notes, this approach requires that the control variables span the Fed’s reaction function, and suffers from the possibility that not all relevant information may be captured by the lower-frequency forecasts. Hansen and McMahon (2016) also study the effects of both types of communication in a low-frequency setting.

<sup>9</sup>Doh et al. (2020) and Handlan (2020) are able to control more flexibly for the Fed staff’s information than Miranda-Agrippino and Ricco (2021) by using machine-learning and text-based techniques applied to the Fed’s alternative policy statements. Cai et al. (2021) and Lakdawala (2019) introduce alternative methods for controlling for the Fed’s *private information* (the difference between the Fed’s and the public’s information) which, as discussed in Section 1.4.2, is necessary for properly “removing” information effects.

<sup>10</sup>This negative correlation does not serve as evidence “against” the presence of information effects—it only rejects the notion that information effects are the only shock operating when the Fed makes announcements. Put differently, without information effects, the negative correlation might be even stronger. Ultimately, the puzzle that suggested the presence of information effects was the effect of interest rate surprises on expectations of real macroeconomic

stantially from the responses of interest rates—e.g., output expectations—instead provides more power for identification.<sup>11</sup> By studying GDP expectations, I directly address one of the main puzzles in the high-frequency literature: that both output and output expectations increase in response to positive interest-rate surprises. Additionally, since my measure is not constructed using financial market data, I eliminate the potential confounding role of risk premia that pervade the “second variables” used by other authors (breakeven inflation and equity returns). As a consequence, my high-frequency measures of GDP and inflation expectations are positively correlated with lower-frequency measures and uncorrelated with measures of inflation risk-premia, in contrast to these alternative second variables. Finally, my identification approach does not require the a priori imposition that the identified shocks have theoretically consistent effects. In contrast to my estimates, this approach only provides set identification.

Why are the estimated effects of (both aspects of) monetary policy so large here? The reduction of measurement error. Because an information shock and monetary shock that raise interest rates have opposite effects on macroeconomic variables, estimated effects that do not attempt to separate the two shocks—or do not do so adequately by, e.g., including controls that not entirely span the Fed’s information set—will exhibit measurement error, and be (at best) attenuated towards zero or even biased in the wrong direction. Approaches that bring in noisy measures of economic fundamentals—e.g., those contaminated by risk premia—will also be biased. Turning away from measurement issues, and towards mechanisms, I show that *Fed-induced* changes in macroeconomic expectations are important determinants of (professional forecasters’) macroeconomic expectations more generally, which suggests an important role for an expectational channel of monetary policy communications.

The paper proceeds as follows. In Section 2, I present my text-based proxies of high-frequency

---

variables (GDP, unemployment, etc.), not stock returns. D’Amico and King (2017) overcome this issue using a set-identification approach. However, they use lower-frequency data on expectations, which may not incorporate all relevant macroeconomic information to which the Fed responds. Andrade and Ferroni (2021) similarly overcome this issue by studying market-based measures of inflation expectations, combined with a sign-restriction approach.

<sup>11</sup>This, combined with the fact that Jarociński and Karadi’s approach only provides set identification of the shocks, allows the authors to (statistically) learn very little about the effects of Fed information—their identified set of impulse responses (that are robust to the chosen prior) includes zero for most horizons.

macroeconomic forecast revisions. In Section 3, I present a simple theoretical framework that explains the identification challenge. In the context of that model, I lay out my identification assumptions and estimate the structural shocks. Section 4 contains evidence on the effects of these shocks on macroeconomic outcomes and expectations, along with a comparison with the effects estimated using existing measures. In Section 5, I dig deeper into both shocks and show that exogenous shocks to overnight interest rates are a thing of the past. Additionally, most of the “information effects” gleaned by markets reflect demand-type factors. Section 6 is the conclusion.

## 1.2 Data: Construction and Validation

A novel aspect of this research is the construction of a high-frequency (HF) proxy for output expectations. Why output? Output is a variable that features in nearly every macroeconomic model, and is the variable whose relationship to traditional HF estimates of monetary shocks is of first-order concern. Why high frequency? Measuring changes in high frequency allows me to isolate the source of the change: Here, the Fed. This ensures that the shocks I estimate are indeed Fed-based shocks and not confounded by information from other sources. Lacking access to existing high-frequency measures of output expectations, however, I use an alternative data source to create my proxy: newspaper articles.

At its core, the text analysis I employ in this paper is essentially a counting exercise: I count co-occurrences of output-related words with words that indicate whether an object is increasing or decreasing. Text-analysis methods for analyzing Fed announcements have used increasingly realistic models of natural language,<sup>12</sup> but with this has come a loss of interpretability and replicability. In contrast, my approach, by design, is conceptually simple, transparent, and easily replicable.

In Section 1.2.1, I describe the construction of my expectations proxy. The simplicity of my construction comes at a cost: It is subject to the criticism of being subjectively designed. I therefore also use a complementary construction that is a bit more complex but removes some subjectivity

---

<sup>12</sup>Examples include the use latent Dirichlet allocation as in chapter 3 and Hansen et al. (2019), whose topics are weighted averages of every word in a corpus; structurally rotated latent semantic analysis (Ter Ellen et al., forthcoming); and more general neural networks (Handlan, 2020).

from the process. My results are nearly unchanged. I then turn, in Section 1.2.2, to an exercise that serves to validate my measure. I show that, in levels and differences, my proxies are positively correlated with existing lower-frequency measures of expectations. Finally, in Section 1.2.4 I briefly describe the traditional numerical data I use throughout the rest of the paper.

### 1.2.1 High-Frequency Text-Based Proxies of Macroeconomic Expectations

The construction of my index proceeds in three steps. I first construct a set of words related to economic output, a set of words that indicate something is increasing, and a set of words that indicate something is decreasing. I then collect a set of newspaper articles written in a one-day window around FOMC announcements, split them into articles written before and after the announcement, and compute a *directionality index* on pre- and post-meeting articles. I then construct my measure of output forecast revisions as the unpredictable component of post-meeting directionality vis-à-vis pre-meeting directionality. I discuss these steps in turn.

**Words Lists** My first set of words refer to economic output. This list is given by

$$\mathcal{Y} \equiv \{\text{economic\_growth, growth, economi, consumer\_spending, output}\}.$$

The origins of this list are the triplet `output, growth, and economy`. In order to show that the results of this paper are robust to expanding or modifying this word list, I trained the popular natural language model of Mikolov et al. (2013) on a large corpus of newspaper articles and extended the initial triplet by extracting—from the model—synonyms of the triplet. The list of synonyms, sorted by their proximity to the triplet, was sensible out to seven words. Thus, in my main analysis I retained the top five most similar words, and later show the robustness of my results to using the top three and top seven most similar words. The details of constructing the list of synonyms are provided in Appendix E.2.

The next two word lists come from the Harvard IV-4 dictionary. The first list contains words that indicate an object is increasing. This set,  $\mathcal{I}$ , consists of all words in the *increase* and *rise*

lists from the Harvard dictionary. The second set,  $\mathcal{D}$ , consists of words indicative of an object decreasing: These are the words from the *decrease* and *fall* word lists. In both cases I retain the unique set of “stemmed” words (the lexical root of words). This ensures that I count all variants of these words in my newspaper articles. For example, *increase*, *increases*, and *increased* are all counted as mentions of *increase*.<sup>13</sup> These word lists are provided in Appendix E.1.

**Newspaper Data** I analyze all newspaper articles written the day before, day of, and day after each FOMC meeting in the *New York Times*, *Wall Street Journal*, and *Washington Post*. I collect these articles from Factiva, searching for articles with the keywords `Federal Reserve` and `FOMC`.<sup>14</sup> The 3-day window ensures that I capture the many articles written on the first day of two-day FOMC meetings (these typically start “The Fed begins a two-day policy meeting today”), and the print articles written on the day after the meeting. I found the timestamps in Factiva to be fairly unreliable, so I manually sorted all articles on the day of the FOMC meeting based on whether they were written before or after the meeting took place. Similarly, I removed duplicate articles.<sup>15</sup>

By considering a 3-day window, I potentially include articles that are not primarily written about the policy meeting. To guard against this, I only retain articles that have `fed` or `rate` in the title or have one of `{fed, fomc, federal}` within five words of one of `{meet, meeting, policy, decision, rate}`. Thus, while the window is not as high frequency as tick-level data, the ability to narrow articles to those that specifically discuss the policy meeting reduces the possibility that other events influence my analysis.<sup>16</sup>

---

<sup>13</sup>I use the stemmer of Porter (1980).

<sup>14</sup>`Fed` is too general, and returns articles related to food and eating.

<sup>15</sup>To do this, I created a vector representation of each article  $i$ , denoted by  $v_i$ , whose length was equal to the total number of unique words across all articles. The  $j$ -th element is the number of times term  $j$  appears in document  $i$ . I then calculate the pairwise distances between all articles using the cosine distance metric: For two documents  $v_i$  and  $v_k$ , this is  $1 - v_i'v_k / [\sqrt{v_i'v_i} + \sqrt{v_k'v_k}]$ . Plotting the distribution of cosine similarities, I found a second mode at 0.97, so I randomly chose one from each set of articles with mutual cosine distance above 0.97.

<sup>16</sup>Upon inspecting the articles, I found that many articles that fit the aforementioned criteria for being “relevant” instead discussed how another country’s central bank might react to the Fed’s announcement. Those articles almost always started with the name of a non-US city (i.e., reflecting where the article was written), so I exclude articles that

Finally, I concatenate all pre- and post-meeting articles into a single document for each meeting  $t$ , given by  $\text{PRE}_t$  and  $\text{POST}_t$ , respectively. I stem all words to their lexical root using the algorithm of Porter (1980). Because my main measure counts co-occurrences of output words and modifiers, I follow Lucca and Trebbi (2009) and break each document into standalone “chunks” of text. This is performed by first tagging each word in a sentence with its part of speech, then using grammatical rules to separate the sentence into smaller constituent sentences (for long or run-on sentences) and various compound phrases. A chunk of text is either a complete sentence with no nested complete sentences or a noun phrase (e.g., `low inflation`) not contained in a complete sentence. The definition of a complete sentence follows rules of the English language and relies on the grammatical parsing of each sentence. I parse sentences using the algorithm of Manning et al. (2014). Though still an imperfect way to determine whether one word modifies another, this method (anecdotally) produces more sensible classifications than determining modification based on proximity (though my results are robust to a proximity-based approach).

**Directionality Index** Having gathered a list of words and documents to analyze, I then put these together to form a directionality index for economic output. For each aggregated document  $d \in \{\text{PRE}_t, \text{POST}_t\}$  with sentence chunks indexed by  $c$ , the index is defined as


$$\omega_d \equiv \underbrace{\frac{\sum_{c \in \mathcal{D}} \sum_{w \in \mathcal{T}_i} \mathbf{1}\{w \in c\} [\sum_{i \in \mathcal{I}} \mathbf{1}\{i \in c\} - \sum_{d \in \mathcal{D}} \mathbf{1}\{d \in c\}]}{\sum_{c \in \mathcal{D}} \sum_{w \in \mathcal{T}_i} \mathbf{1}\{w \in c\} [\sum_{i \in \mathcal{I}} \mathbf{1}\{i \in c\} + \sum_{d \in \mathcal{D}} \mathbf{1}\{d \in c\}]}}_{\text{Sentiment Score}} \times \underbrace{\frac{1}{|d|}}_{\text{Normalization}}. \quad (1.1)$$

The “sentiment score” component of this measure is so named because it resembles sentiment scores found across the natural language processing literature, whose typical construction is the percent difference in counts of “positive” and “negative” words. Here, the score is the number of times a word from the topic is modified by an increasing word less the number of times a word from the topic is modified by a decreasing word, divided by mentions in either direction. For chunks that contain a negation (`not` and `n't`), I flip the sign of the bracketed term in the numerator. Dividing

---

have one of these non-US city names within the first 5% of the text. The list of these cities is in Appendix E.1.

Figure 1.1: Newspaper-based Output Directionality and Macroeconomic Expectations

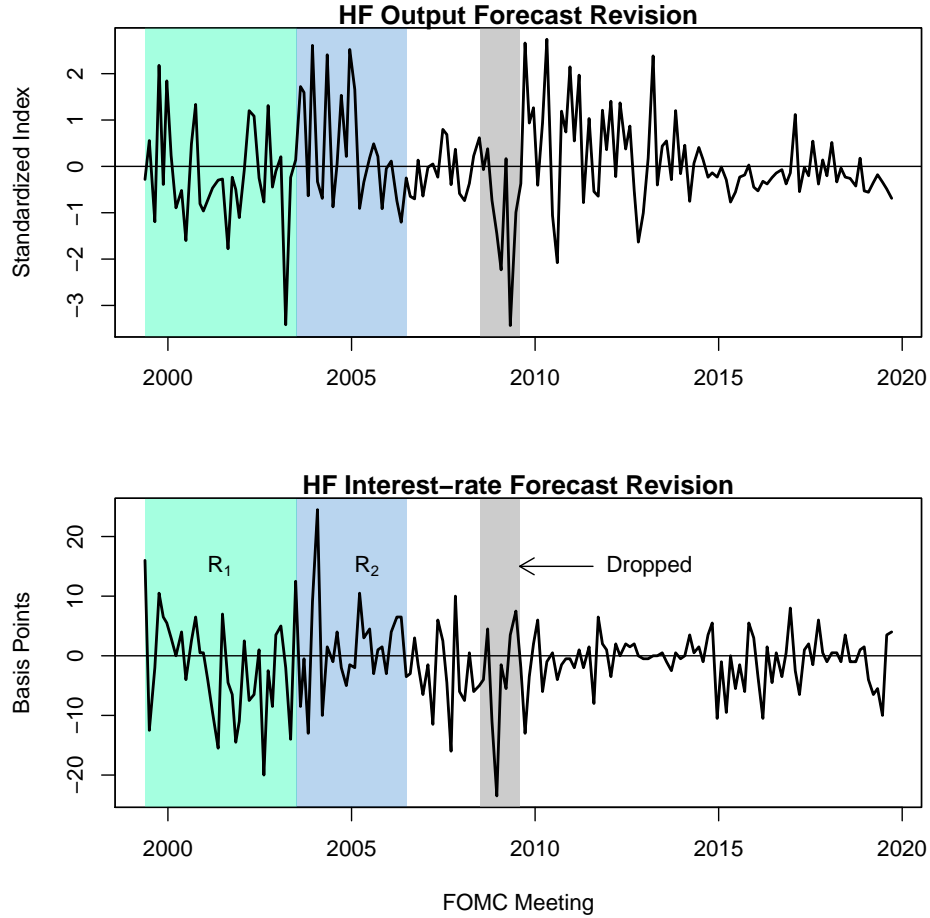


The left panel plots the average GDP forecast from the Blue Chip (black dashed line) described by equation (1.3) against the 1-year (eight-meeting) moving average of post-meeting output directionality  $\omega_{POST_t}$  (solid red line). The right panel shows a scatter plot of the two series (in which each meeting month is merged with the corresponding Blue Chip month), using only observations that are used in the baseline empirical sample (excluding the July 2008–July 2009 period and meetings that occur in the first 7 days of each month). For ease of comparison, I normalize  $\omega_{POST_t}$  to have the same mean and standard deviation as the corresponding  $\Upsilon_{t|t}$  variable.

by the total number of sentences (the “normalization”) gives the index the average per-sentence (chunk) sentiment score.

In levels, I posit the economic growth directionality index as a proxy for macroeconomic expectations. The time series of  $\omega_{POST_t}$  in the left panel of figure 1.1 (which is discussed in more detail in Section 1.2.2) supports this interpretation. In Appendix E.3, I create an alternative directionality index that makes this link explicit. Briefly, I count all pairwise co-occurrences of words in slightly expanded versions of  $\mathcal{Y}$ ,  $\mathcal{D}$ , and  $\mathcal{I}$  in a large corpus of newspaper articles. I estimate a LASSO regression to predict the level of real GDP expectations from the Blue Chip survey using these co-occurrence counts. This gives me a mapping from words to GDP expectations. I apply this mapping to pre- and post-FOMC articles to generate my alternative directionality indexes. The co-occurrences “selected” by LASSO are sensible. However, because my results are essentially unchanged between my baseline measure and this alternative, I opt to use the simpler index presented in equation (1.1).

Figure 1.2: High-frequency Forecast Revisions of Output and Interest Rates




The top panel plots the high-frequency output forecast revision measure (from newspapers) described in Section 1.2.1 (the residual from equation 1.2). The bottom panel shows the 30-minute change in the price of the fourth-quarter Eurodollar futures, described in Section 1.2.4. The first two shaded areas present the regimes used for identification, described in Section 1.3.2. The third, gray shaded area highlights observations that are dropped for most analyses, discussed in Section 1.2.4.

**Proxy for Changes in Macroeconomic Expectations** The last step is to take the directionality indexes from levels to differences. One complication arises in this step: Namely, the indexes constructed above do not distinguish between descriptions of the Fed’s announcement (i.e., the “level” of the announcement) and the surprise component of that announcement. I therefore construct my proxy for changes in output expectations as the unpredictable component of post-meeting directionality vis-à-vis pre-meeting directionality. Formally, I construct the proxy  $\hat{y}_t$  as the residual from the following regression:

$$\omega_{\text{POST}_t} = a + b \omega_{\text{PRE}_t} + c \omega_{\text{POST}_{t-1}} + \hat{y}_t. \quad (1.2)$$

Figure 1.3: Blue Chip GDP Forecast Revisions and the High-frequency Proxy



The left panel plots the average Blue Chip forecast revision over the next three forecast horizons ( $\bar{Y}_{t|t+1} - \bar{Y}_{t|t}$  from equation (1.3)) against my HF measure of output forecast revisions,  $\hat{y}_t$ , standardized to have unit variance over the plotted sample. The slope of the line thus corresponds to the estimate of  $\beta$  from equation (1.4) in Table 1.1. The right panel shows this estimated slope using different Blue Chip forecast horizons—from zero to four quarters ahead. In that panel the right- and left-hand-side variables are standardized, so the coefficient is a correlation coefficient. The sample is as described in Section 1.2.4. Confidence intervals are generated using robust standard errors.

The constant and slope coefficients in the regression are 0.0002, 0.18, and 0.11.<sup>17</sup> This suggests that a fair amount of post-meeting articles discuss the surprise component of the policy announcement; to quantify this differently, the  $R^2$  from the regression is only 0.13, which suggests that a fair amount of post-meeting coverage is unpredictable based on pre-meeting coverage. Without the inclusion of  $\omega_{\text{POST}_{t-1}}$ , the index exhibits a small degree of autocorrelation. Otherwise its inclusion is inconsequential for my estimates. The top panel of Figure 1.2 shows the time series of  $\hat{y}_t$ .

## 1.2.2 Validation

This section contains several exercises meant to validate the indexes constructed in Section 1.2.1 as proxies for macroeconomic expectations. I show, first, that the directionality indexes (the  $\omega_t$  variables) are positively correlated with the level of real GDP expectations from the Blue Chip

<sup>17</sup>The heteroskedasticity-robust t-statistics are 4.9, 3.6, and 1.5, respectively. I estimate this regression over my full sample of newspapers, which starts in 1995, but the regression coefficients are almost identical using the 1999+ sample (respectively, 0.0002, 0.18, and 0.19).

survey of the same month. I then show that the high-frequency forecast revisions I construct,  $\hat{y}_t$ , are positively correlated with lower-frequency forecast revisions. Finally, I perform a case study to show that the indexes are picking up features of the text that are noticeable to a human reader.

Figure 1.1 shows that the macroeconomic directionality indexes are positively correlated with macroeconomic expectations taken from the Blue Chip survey. The dashed black line in the left panel is a summary statistic of GDP expectations I use often throughout the paper:

$$\mathbb{Y}_{t|\tau} \equiv \frac{1}{3} \sum_{h=1}^3 \mathbb{E}_{\tau}^{\text{Blue Chip}} [\Delta \text{Real GDP}_{t+h}]. \quad (1.3)$$

This summary statistic—used by, e.g., Nakamura and Steinsson (2018)—is the average forecast of real GDP growth made in month  $\tau$  over the year starting in month  $t$  (in Figure 1.1, I set  $\tau = t$ ).<sup>18</sup> The solid blue line in the left panel is the 1-year (eight-meeting) moving average of post-meeting directionality  $\omega_{\text{POST}_t}$ .<sup>19</sup> The lines are clearly positively correlated, and the right panel confirms that this correlation is not the result of one or two influential observations. The points in the right panel plot  $\mathbb{Y}_{t|t}$  against  $\omega_{\text{POST}_t}$  (not the moving average) for the months in my baseline sample.<sup>20</sup> The correlation of these series is 0.36, with a robust standard error of 0.08.<sup>21</sup> In summary, my output expectations are correlated with traditional measures of output expectations, in levels.

My index of output forecast *revisions* is also correlated with traditional measures of output expectation revisions. Figure 1.3 shows that the high-frequency proxy for macroeconomic forecast revisions,  $\hat{y}_t$ , is also positively correlated with lower-frequency forecast revisions taken from the Blue Chip survey. The left panel plots the 1-month Blue Chip forecast revision surrounding each FOMC meeting ( $\mathbb{Y}_{t|t+1} - \mathbb{Y}_{t|t}$ ) against the high-frequency proxy. This positive correlation is

---

<sup>18</sup>Formally, the average forecast of GDP growth for quarters  $q(t)$ ,  $q(t) + 1$ , and  $q(t) + 2$ , where  $q(t)$  is the quarter of month  $t$ .

<sup>19</sup>The graph with pre-meeting directionality looks very similar.

<sup>20</sup>I exclude the observations from July 2008 through July 2009, which will be removed from my sample later in the paper because of the asset-pricing anomalies over this period discussed by Nakamura and Steinsson (2018). Also, to make the sample consistent with my later analysis, I exclude months with FOMC meetings that occur in the first 7 days of the month since, as Nakamura and Steinsson (2018) note, the exact timing of when Blue Chip respondents complete their surveys is not clear.

<sup>21</sup>The correlation with pre-meeting directionality ( $\omega_{\text{PRE}_t}$ ) is 0.37 (s.e. 0.08).

Table 1.1: Low- and High-frequency Forecast Revisions: Robustness

$\hat{y}$ (HF newspaper)	0.0267 (2.72)		0.0258 (2.68)		0.0191 (1.95)
Interest-rate surprise		0.00452 (3.17)	0.00439 (3.20)	0.00482 (2.85)	0.00207 (1.71)
HF S&P 500 return				0.0574 (2.42)	
Jobs number					-0.00896 (-7.30)
Jobs number, surprise					-0.0546 (-0.50)
One-month S&P Return					0.00696 (3.65)
Lagged BBK index					0.0410 (3.82)
Observations	131	131	131	138	131
$R^2$	0.0566	0.0629	0.116	0.0684	0.373
$\overline{ \text{LHS} }$	0.0775	0.0775	0.0775	0.0775	0.0775

$t$ -statistics computed using robust standard errors are in parentheses, since the point estimates are so small.

This table shows estimates of equation (1.4). The left-hand-side is the 1-month forecast revision of GDP growth over the next three forecasting horizons, presented in equation (1.3), from the Blue Chip survey, in percentage points. The HF output forecast revision  $\hat{y}_t$  is from newspapers and standardized over the regression sample. The interest-rate measure ( $\hat{i}_t$ ) is the change in the four-quarter Eurodollar future contract, in basis points. The 1-month stock return is the 4-week percentage return in the S&P 500 (ending 1-day before the FOMC announcement). The lagged BBK index is the standardized value of the index of Brave et al. (2019) from one-month before the date  $t$  FOMC meeting. “Jobs number” and “Jobs number, surprise” are the level and surprise component of month  $t$ ’s release of the change in non-farm payrolls (in units of 100,000 jobs). Expectations of that release are from Bloomberg. The sample consists of all regularly scheduled FOMC meetings between May 1999 and October 2019, excluding July 2008–July 2009, that occur after the first week of the month. The row  $\overline{|\text{LHS}|}$  is the average absolute GDP forecast revision over the regression sample (0.3 percentage point).

encouraging. I put this correlation to the test by estimating the regression

$$\mathbb{Y}_{t|t+1} - \mathbb{Y}_{t|t} = \alpha + \beta \hat{y}_t + e_t, \quad (1.4)$$

where  $X_t$  contains various explanatory variables of interest. The results are in Table 1.1. The first column shows that the correlation is indeed statistically significant—a finding that remains consistent across the columns. The correlation is not due to the construction of the Blue Chip

variable in equation 1.4; the right panel breaks down the correlation by Blue Chip forecast horizon. This shows that  $\hat{y}_t$  is fairly evenly correlated with forecast revisions out to about 1 year.

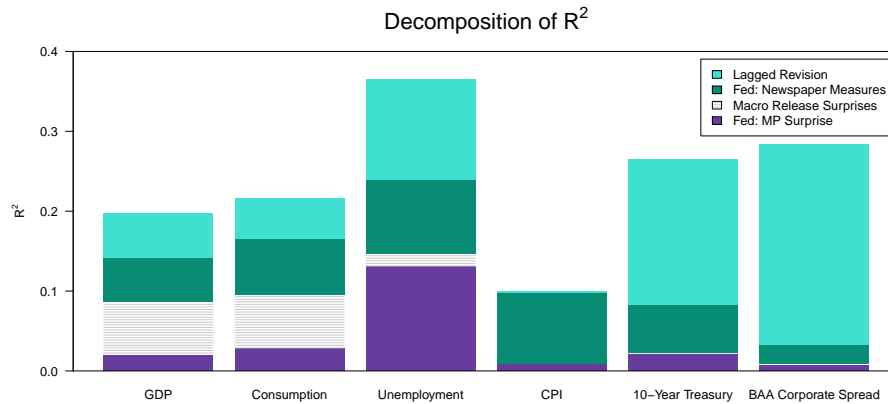
Table 1.1 also highlights why  $\hat{y}_t$  is a useful addition to the study of monetary policy announcements. In the second column, I only include a traditional estimate of monetary shocks on the right-hand-side of equation (1.4)—HF changes in 1-year interest-rate expectations, denoted by  $\hat{i}_t$  (displayed in Figure 1.2 and described in Section 1.2.4). This column is a reproduction of a known, but puzzling, finding: Surprise interest-rate increases (i.e., a contractionary traditional HF shock) cause GDP expectations to increase. Interestingly, when the HF output revision is added to that regression, the coefficient on the interest-rate surprise remains nearly unchanged. This is a hint that  $\hat{y}_t$  is not a useful variable for controlling for information effects—if it were, the coefficient on the interest-rate surprise should change. This also highlights (thinking about the omitted-variables-bias equation) that  $\hat{y}_t$  and  $\hat{i}_t$  are only weakly positively correlated (the correlation coefficient is 0.09, with a  $t$ -statistic of 1.3). The positive correlation suggests that information effects are still present in high frequency. The fact that the correlation is fairly weak also suggests a useful feature of bringing  $\hat{y}_t$  to this identification problem: It contains independent variation. This is less true with stock prices (Jarociński and Karadi, 2020); the high-frequency S&P 500 return around Fed announcements with  $\hat{i}_t$  is  $-0.36$ , with a  $t$ -statistic of  $-3.9$ .<sup>22</sup>

In the last column, I show that  $\hat{y}_t$  continues to be positively correlated with lower-frequency expectations when controlling for recent macroeconomic news. Bauer and Swanson (2020) estimate similar specifications using low-frequency measures of GDP forecast revisions. They find that recent news predicts both interest rate-surprises and GDP forecast revisions, and substantially mitigates the positive correlation between the two. They interpret these results as a suggestion that recent macroeconomic news induces a spurious positive correlation between interest-rate surprises and GDP forecast revisions, which argues against the presence of information effects. Following

---

<sup>22</sup>One reaction to this negative correlation would be to ask “doesn’t this refute the presence information effects?” It does not. This negative correlation is only useful for rejecting that there are *only* information effects (or that they are the strongest determinant of stock prices). Put differently, without information effects, this correlation would be much more negative.

Figure 1.4: Shapley Decomposition of Forecast Revision  $R^2$



Bauer and Swanson, I include the most recent level of the non-farm payrolls release, the surprise component of that release, the 1-month stock return, and the lagged index of Brave et al. (2019). The coefficient drops somewhat, but remains marginally statistically significant. For transparency, this coefficient drops nearly to zero if the 13-week return of the S&P 500 return is included instead (the exact variable used by Bauer and Swanson). However, in Appendix I, I show that the 13-week return is the most potent return horizon over which stock returns “predict” news around monetary policy announcements, which calls into question the robustness of the conclusions of Bauer and Swanson

While I have largely discussed the positive correlation between  $\hat{y}_t$  and the lower-frequency measure of forecast revisions as a reassuring feature of the construction of  $\hat{y}_t$ , it is worth noting that this correlation is not a necessary condition for  $\hat{y}_t$  to be a good measure of GDP expectations. Recall that  $\hat{y}_t$  is a measure of *Fed-induced* changes in expectations about GDP. If the Fed’s news tended to be dwarfed by other events that occurs over the course of a typical month, this correlation would tend to zero. So, the finding of a positive correlation suggests that the Fed plays some non-trivial role in shaping GDP expectations. To push this point a bit further, in figure 1.4 I perform a Shapley decomposition, of the  $R^2$  of regressions of one-month changes Blue Chip expectations on my newspaper-based measures of GDP and inflation forecast revisions; a high-frequency measure of interest rate forecast revisions; the surprise component of GDP, CPI inflation, unemployment,

and the non-farm payrolls releases in each month (described below); and the lagged forecast revision (to capture information frictions). I decompose the  $R^2$  according to these four groups of variables. As described by Shorrocks (1982), this Shapley decomposition decomposes the  $R^2$  in a regression and takes into account the potential correlations among variables in the regression. The Fed-induced changes in macroeconomic expectations (the dark solid bars) contribute as much, if not more, to the determination of revisions as macroeconomic announcement surprises. This is further evidence that the Fed plays an important role in determining macroeconomic expectations. The quantitative magnitude of these effects are larger than would be suggested by defining policy narrowly using only interest rate forecast revisions, as typically considered in the literature.

A benefit to performing this “automated narrative analysis” is that I can read newspaper articles to corroborate the largest shocks. For example, the largest negative Fed-induced change in expectations about GDP came in March 2004, when Greg Ip’s post-meeting WSJ article state:

The slightly less upbeat tone of the statement drove long-term bond yields down sharply. The Fed said risks to economic growth remain “roughly equal” while the risk of an “unwelcome fall in inflation” was “almost equal” to that of a rise in inflation.

In contrast, pre-meeting WSJ coverage states that yields were not “likely to fall much further, given countercurrents of strong economic growth” and that the Fed was not expected to “tweak significantly the language in the accompanying policy statement.”

### 1.2.3 Predictability

When proposing an estimate of a macroeconomic shock, it is important to examine the extent to which that shock is predictable in some way. I test the predictability of my newly constructed estimate of output forecast revisions,  $\hat{y}_t$ , and a traditional measure of monetary shocks,  $\hat{i}_t$  (described in Section (1.2.4)) in Table 1.2. Specifically, I regress each measure on a slew of controls. In the first two columns, I include two lags of the left-hand-side variables, the most recent surprise component of three macroeconomic news releases, and the two most recent changes in non-farm payrolls (using real-time data). All of these variables are observable to markets and Fed watchers before

Table 1.2: Predictability of High-frequency Output ( $\hat{y}_t$ ) and Interest-rate ( $\hat{i}_t$ ) Forecast Revisions

	$\hat{y}_t$	$\hat{i}_t$	$\hat{y}_t$	$\hat{i}_t$	$\hat{y}_t$	$\hat{i}_t$
Lagged LHS	0.035 (0.082)	-0.049 (0.093)	-0.008 (0.080)	-0.105 (0.087)	0.031 (0.083)	-0.055 (0.092)
Twice-Lagged LHS	0.086 (0.095)	-0.093 (0.096)	0.071 (0.097)	-0.125 (0.099)	0.098 (0.098)	-0.099 (0.098)
BB Surprise: CPI	0.052 (0.055)	-0.094 (0.384)	0.052 (0.053)	-0.077 (0.310)	0.052 (0.055)	-0.093 (0.372)
BB Surprise: GDP	0.072 (0.041)	-0.160 (0.285)	0.058 (0.039)	-0.257 (0.266)	0.070 (0.042)	-0.179 (0.282)
BB Surprise: Jobs	-0.473 (1.574)	-5.870 (11.600)	0.077 (1.703)	-4.321 (10.676)	-0.401 (1.576)	-5.577 (11.603)
First $\Delta$ NFPR	1.905 (1.372)	9.573 (7.180)	0.358 (1.570)	2.702 (6.528)	1.794 (1.399)	8.943 (6.975)
Second $\Delta$ NFPR	-1.181 (1.136)	-0.554 (5.452)	-1.323 (1.197)	-0.013 (5.321)	-1.102 (1.142)	0.056 (5.365)
1-quarter stock ret.			0.016 (0.014)	0.188 (0.099)		
BBK Index $_{m(t)-1}$			0.377 (0.158)	1.259 (0.972)		
6-week stock ret.					0.018 (0.018)	0.111 (0.136)
Constant	-0.039 (0.114)	-2.007 (0.737)	0.086 (0.136)	-1.806 (0.729)	-0.049 (0.113)	-2.106 (0.744)
Observations	159	159	159	159	159	159
Adjusted $R^2$	0.032	0.012	0.108	0.078	0.032	0.012
F	1.277	1.056	2.453	2.008	1.273	1.146
p(F)	0.265	0.395	0.012	0.042	0.262	0.336

This table presents results from regressing the HF measure of output and interest-rate forecast revisions around FOMC announcements on several variables. The output measure ( $\hat{y}_t$ ) is from newspapers and standardized over the regression sample. The interest-rate measure ( $\hat{i}_t$ ) is the change in the 4-quarter Eurodollar future contract, in basis points. In each column two lags (at FOMC meeting frequency) of the left-hand-side variables are included. The three “BB surprise” variables are the difference between expected and announced annualized CPI inflation (basis points), annualized GDP growth (basis points), and monthly change in non-farm payrolls (100,000 jobs). I compute these forecast errors (from Bloomberg) from the most recent (pre-FOMC) announcement of each variable. The two NFPR variables are the most recent (pre-FOMC) and second most recent real-time releases of non-farm payrolls (100,000 jobs). The 1-quarter stock return is the 13-week percent return in the S&P 500 (ending one day before the FOMC announcement). The lagged BBK index is the standardized value of the index of Brave et al. (2019) from 1-month before the date  $t$  FOMC meeting. The sample consists of all regularly scheduled FOMC meetings between May 1999 and October 2019. Robust standard errors are in parentheses. The row labeled “F” is the F-statistic from a joint test that all coefficients (except the constant) are nonzero, and p(F) gives the associated p-value.

the meeting announcement. None of these right-hand-side variables have a statistically significant coefficient. Jointly the variables are statistically insignificant. The  $R^2$  of both regressions are very small. The finding of minimal autocorrelation is consistent with the findings of Miranda-Agrippino and Ricco (2021).

In the remaining columns I add additional variables considered by Bauer and Swanson.<sup>23</sup> Following the authors, in the third and fourth columns I add the 13-week return of the S&P 500 and the 1-month lagged index of Brave et al. (2019) (“the BBK index”). Here the evidence of predictability is stronger, with joint tests of significance smaller than 5% for both variables. In the fifth and sixth columns I examine these findings further. First, noting that the BBK index is not observable, I exclude it from the regression.<sup>24</sup> Second, recalling, from the discussion of table 1.1, the sensitivity of “predictability” regressions to the stock-market horizon, I change the return horizon to a six-week horizon. This approximates the inter-meeting stock return, since eight scheduled meetings are held per calendar year. The results in these columns again suggest a lack of predictability, again calling into question the robustness of the conclusion Bauer and Swanson, that both GDP and interest-rate forecast revisions are predictable from publicly available information.

#### 1.2.4 Numerical Data

To measure the surprise component of monetary policy decisions, I use tick-level data on Federal Funds and Eurodollar futures. I use the 30-minute change in the current-month Federal Funds future rate (in the spirit of Kuttner (2001)) and the change in the price of the 4-quarter-ahead Eurodollar futures contract.<sup>25</sup> These series are well-known in empirical monetary economics and were extended by Acosta and Saia (In progress) through 2019 using tick-level data purchased from the CME group. The authors followed Nakamura and Steinsson (2018) exactly in the construction

---

<sup>23</sup>The authors also include recent changes in non-farm payrolls. Using real-time data, I find that those variables are not predictive of monetary shocks.

<sup>24</sup>The index did not exist for most of my sample. Even when it does exist, the index for month  $m(t)$  (used in the regression) is not until month  $m(t) + 2$ . Thus, when considering an FOMC meeting at time  $t$ , the value of the index for month  $m(t) - 1$  is not available to markets.

<sup>25</sup>For those familiar with the variable names in HF papers, this is ED4.

of these variables. The 4-quarter Eurodollar is highly (0.95) correlated with the shock used by Nakamura and Steinsson and the path factor of Gürkaynak et al. (2005), with a correlation coefficient of 0.91. Here, I measure high-frequency S&P 500 prices using the exchange traded fund SPY, and low-frequency prices from Yahoo finance.

I take data on macroeconomic expectations from Blue Chip Economic Indicators. In some regression specifications I control for the surprise components of macroeconomic news releases. I collect expectations and the announced values of these variables from Bloomberg—the data appendix of Acosta and Saia (In progress) describes these data in detail. When studying macroeconomic effects I use the shadow Federal Funds rate of Wu and Xia (2016) and other macroeconomic aggregates from FRED.

Unless noted otherwise, my sample consists of all regularly scheduled FOMC meetings between May 1999 and October 2019. May 1999 is the first meeting after which the Fed started to regularly release post-meeting statements. October 2019 is when my high-frequency data end. I exclude observations from July 2008 through 2009 because of the asset-pricing anomalies over this period discussed by Nakamura and Steinsson (2018).

### **1.3 Shock Identification**

In Section 1.3.1, I describe an illustrative version of the model I use to estimate my shocks, which models how expectations are revised in light of Fed announcements. This allows me to formalize the identification challenge: Forecast revisions about observable variables (e.g., interest rates) are linear combinations of forecast revisions about structural shocks (e.g., monetary and information shocks). This implies that no high-frequency change in expectations or prices will identify macroeconomic shocks *unless* additional assumptions are made. As a consequence, my new measure of high-frequency output expectations is neither an estimate of “information effects” nor can it be used as an instrument for information effects.

This is not to say that all hope is lost; only that more work has to be done. The second purpose of my illustrative model is to show that, together, my new measure, traditional estimates of mone-

tary shocks, and historical information about the nature of the Fed’s communication practices can be used to separately identify monetary and information shocks using the heteroskedasticity-based identification assumptions of Rigobon (2003). I pair the illustrative model with my data to show the intuition behind this procedure in Section 1.3.2. Finally, in Section 1.3.3 I describe the formal model and identification assumptions, then present the results of the estimation in Section 1.3.4.

### 1.3.1 The Identification Challenge in an Illustrative Model

My illustrative model describes how markets (and Fed watchers more generally) form macroeconomic expectations, which allows me to study how these expectations are revised in response to the Fed’s policy announcements. Fed watchers use a simple macroeconomic model for this task. While this model, in its simplicity, serves its expositional purpose, when appropriate I discuss why the intuition gleaned from the simple model is consistent with expanded elaborations of the model.

**Economy** The macroeconomic model used by forecasters is described by output  $y_\tau$  and the interest rate  $i_\tau$ , which are related by the following two equations:

$$i_\tau = \phi y_\tau + \varepsilon_\tau \tag{1.5a}$$

$$y_\tau = -\gamma i_\tau + \eta_\tau. \tag{1.5b}$$

The exogenous shocks are a monetary shock,  $\varepsilon_\tau$ , and a macroeconomic fundamental,  $\eta_\tau$ . These serially uncorrelated shocks are normally distributed with zero mean and variances  $\sigma_\varepsilon^2$  and  $\sigma_\eta^2$ . For the purposes of building intuition, I assume that the coefficients  $\phi$  and  $\gamma$  are positive. I do not use these sign restrictions in my identification procedure. In appendix A, I show that this model nests the textbook three-equation New Keynesian model (Galí, 2015), in which I substitute out inflation and remove technology shocks (thus, the information shock here is a demand shock in the context of the New Keynesian model). As such, I refer to the equations as the forecaster’s Taylor Rule (Taylor, 1993) and IS equation, respectively. The elimination of supply shocks is consistent with my findings below; whereby monetary announcements reveal very little about supply shocks.

Appendix C highlights the fact that the intuition from this static bivariate model extends to linear forward- and backward-looking general equilibrium models with noisy information as in, e.g., Blanchard et al. (2013).

**Fed Announcements** At discrete points in time, indexed by  $t$ , the Fed makes policy announcements. Empirically, I measure expectations about the observable variables,  $y_\tau$  and  $i_\tau$ , shortly before and shortly after each policy announcement. Denote the time of those pre- and post-announcement measurements by  $\underline{t}$  and  $\bar{t}$ , respectively, and expectations taken at each time  $t$  about  $x$  at time  $\tau$  as  $x_{\tau,t} = \mathbb{E}_t[x_\tau]$ , for  $x \in \{y, i, \eta, \varepsilon\}$ . Equations (1.5a) and (1.5b) show how these expectations of observable variables relate to perceptions of structural shocks before (left column) and after (right column) each announcement

$$i_{\tau,\underline{t}} = \phi y_{\tau,\underline{t}} + \varepsilon_{\tau,\underline{t}} \qquad i_{\tau,\bar{t}} = \phi y_{\tau,\bar{t}} + \varepsilon_{\tau,\bar{t}} \qquad (1.6a)$$

$$y_{\tau,\underline{t}} = -\gamma i_{\tau,\underline{t}} + \eta_{\tau,\underline{t}} \qquad y_{\tau,\bar{t}} = -\gamma i_{\tau,\bar{t}} + \eta_{\tau,\bar{t}}. \qquad (1.6b)$$

Denote high-frequency changes in expectations about  $x_\tau \in \{y_\tau, i_\tau, \eta_\tau, \varepsilon_\tau\}$  by  $\hat{x}_{\tau,t} = \mathbb{E}_{\bar{t}}[x_\tau] - \mathbb{E}_{\underline{t}}[x_\tau]$ . Then taking the difference between the left and right columns of equation (1.6a) reveals

$$\hat{i}_{\tau,t} = \phi \hat{y}_{\tau,t} + \hat{\varepsilon}_{\tau,t}. \qquad (1.7a)$$

Similarly, the IS equation in expectation-revision space is

$$\hat{y}_{\tau,t} = -\gamma \hat{i}_{\tau,t} + \hat{\eta}_{\tau,t}. \qquad (1.7b)$$

Combining equations (1.7a) and (1.7b) allows the observable variables to be expressed as linear combinations of the forecast revisions about perceived structural shocks:

$$\hat{i}_{\tau,t} = \left( \frac{\phi}{1 + \gamma\phi} \right) \hat{\eta}_{\tau,t} + \left( \frac{1}{1 + \gamma\phi} \right) \hat{\varepsilon}_{\tau,t} \quad (1.8a)$$

$$\hat{y}_{\tau,t} = \left( \frac{1}{1 + \gamma\phi} \right) \hat{\eta}_{\tau,t} - \left( \frac{\gamma\phi}{1 + \gamma\phi} \right) \hat{\varepsilon}_{\tau,t}. \quad (1.8b)$$

I refer to  $\hat{\varepsilon}_{\tau,t}$  as a perceived monetary shock, and  $\hat{\eta}_{\tau,t}$  as an information shock or information effects.

**The Identification Problem** Equations (1.8a) and (1.8b) are the crux of the identification challenge. First, equation (1.8a) shows that in general, traditional HF measures of monetary shocks ( $\hat{i}_{\tau,t}$ ) are contaminated by information effects ( $\hat{\eta}_{\tau,t}$ ). There is only a special case—when the Fed has no independent knowledge about the state of the economy (so that  $\hat{\eta}_{\tau,t} = 0, \forall t$ )—in which traditional measures identify the (perceived) monetary policy shock  $\hat{\varepsilon}_{\tau,t}$ . A testable implication of the model is that if traditional estimates are not contaminated by information effects, then the correlation between interest-rate and output expectations must be negative. Empirical evidence refutes this implication. Campbell et al. (2012) and Nakamura and Steinsson (2018) find a positive correlation between output and interest-rate forecast revisions using low-frequency data, which I confirm using high-frequency data.

Second, equation (1.8b) shows that an estimate of output forecast revisions can neither be used to control for information effects nor to instrument for information effects. On the first point, in Appendix D I show that the residual from a regression of interest rate expectations on output expectations (i.e. a “cleaned” interest rate surprise) only identifies monetary shocks in the case that output expectations do not respond to monetary shocks which. This, in turn, is only the case with full monetary neutrality (i.e.  $\phi = 0$ ) or there are no information effects in the first place.<sup>26</sup> (Interestingly, in the absence of information effects, any asset-price change in the window around a

---

<sup>26</sup>In the language of Angrist and Pischke (2008),  $\hat{y}_t$  is a “bad control” for information effects, since it is affected by monetary shocks.

policy announcement identifies a monetary policy shock!) Next, output forecast revisions are not a valid instrument for information effects, because they are not exogenous with respect to monetary policy shocks.

These observations intuitively show the necessity for additional identification assumptions to be made in order to separately identify monetary and information shocks. Let  $(\hat{\sigma}_v, \hat{\sigma}_\eta, \hat{\rho}_{\eta,v})$  be the variance of  $\hat{\eta}_{\tau,t}$ , the variance of  $\hat{\varepsilon}_{\tau,t}$ , and their covariance, respectively. Formally, the identification problem is that the model contains five parameters  $(\phi, \gamma, \hat{\sigma}_v, \hat{\sigma}_\eta, \hat{\rho}_{\eta,v})$ , but the data only provide three empirical moments: the variances of  $\hat{y}_t$  and  $\hat{i}_t$  and their covariance. I next discuss my identifying assumptions in the context of my data and illustrative model. These assumptions will allow me to estimate the linear mapping (a function of  $\gamma$  and  $\phi$ ) between the (unobserved) perceived structural shocks  $(\hat{\eta}_{\tau,t}$  and  $\hat{\varepsilon}_{\tau,t})$  and the (observed) forecast revisions  $(\hat{y}_{\tau,t}$  and  $\hat{i}_{\tau,t})$ . I will then invert that linear mapping to recover the shocks from the revisions.

To this point I have remained agnostic as to the exact mathematical specification of the information communicated by Fed policy announcements. Such a formal structure is not necessary for my identification approach. However, I do make two assumptions about the information conveyed and how markets make use of that information. First, I assume that the information revealed allows forecast revisions about the structural shocks to be uncorrelated, i.e.,  $\hat{\rho}_{\eta,v} = 0$ . This need not be the case in general signal-extraction problems: If the Fed's communications are not sufficiently detailed in their discussion of each type of shock (economic fundamentals,  $\eta_t$  vs. monetary policy,  $\varepsilon_t$ ), then markets will generally have to use their prior knowledge to parse the independent information revealed about each shock. Empirically, this assumption can be tested with overidentification tests. Theoretically, in Appendix B.1, I discuss a particular information structure under which this assumption would be valid. The second assumption is that markets make larger forecast revisions about a variable when signals about that variable are clearer—this is an implication, for example, of Bayes' rule. These assumptions, plus some historical knowledge described in the next section, allow me to identify the model's parameters.

### 1.3.2 Regimes and Intuition for the Identification Assumptions

I bring historical knowledge about the nature of the Fed’s communication practices to the identification problem. Intuitively, this knowledge provides the fourth moment needed to identify the model’s four parameters. To be concrete, I rely on the episode discussed by Lunsford (2020): the August 2003 introduction of interest-rate forward guidance. To give context, in June 2003 the FOMC had lowered the Federal Funds rate to 1%. In August 2003, the post-meeting statement declared that this “policy accommodation can be maintained for a considerable period.” This was the first instance of explicit forward guidance regarding the path of interest rates, and was used for reasons similar to what prompted its major re-emergence in 2008. In 2003, 1% was essentially seen as the effective lower bound on nominal interest rates. This episode thus gives me a natural place to split my sample. My first regime,  $\mathcal{R}_1$ , consists of all meetings from the start of my sample (May 1999) through June 2003. The second regime,  $\mathcal{R}_2$ , extends from August 2003 through the end of 2006.<sup>27</sup> These regimes are shown by the two left-most shaded regions of Figure 1.2.

To see how this episode can provide useful variation for identification, consider the change in the behavior of forecast revisions following the announcement, shown in Figure 1.5. Focusing on panel A, the behavior of forecast revisions clearly changed between the two regimes. In the first regime (left-most plot), GDP and interest-rate expectations tended to be revised in the same direction following Fed announcements. This statistically significant positive correlation disappears—and becomes even slightly negative—following the regime change.


The change in the correlation of GDP and interest-rate forecast revisions is predicted by the illustrative model, given the nature of this episode. In the context of that model, perceived information shocks induce a positive correlation between GDP and interest-rate forecast revisions, while monetary shocks induce a negative correlation. This is evident from equations (1.8a) and (1.8b). Since both shocks are potentially present in both regimes, however, the correlation between the

---

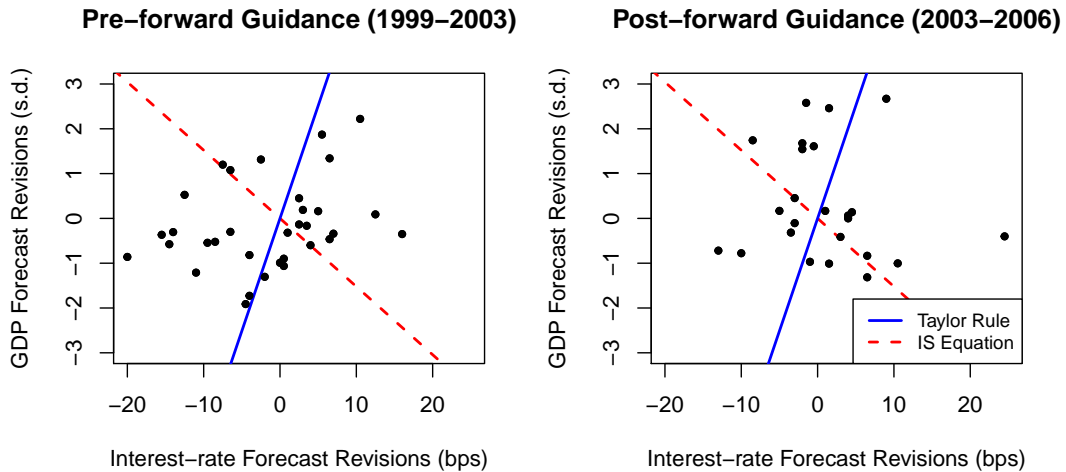
<sup>27</sup>These are slight elongations of the sample period studied by Lunsford (2020). I elongate the sample for statistical precision—my point estimates are nearly identical using Lunsford’s exact regimes, but mildly less-precisely estimated. I will also note that I have estimated an over-identified system using the rest of my sample as a third regime, but again the point estimates are very similar.

Figure 1.5: The Identification Assumptions in Pictures

Panel A: Correlations in the two Regimes



Panel B: Estimated Data-Generating Process



Panels A and B contain scatter plots of high-frequency GDP and interest-rate forecast revisions around Fed announcements ( $\hat{y}_t$  and  $\hat{i}_t$  in the text). The left panel shows observations between May 1999 and June 2003; the right panel features observations between August 2003 and December 2006. The green line in panel A is the unconditional line of best fit, estimated separately in both regimes. The lines in panel B are those estimated to generate the data in both regimes (in a maximum-likelihood sense), estimated using data from both regimes using the identification procedure described in Section 1.3.4.

two series depends on which shock is larger, as determined by their variances. Lunsford (2020) highlights the fact that before 2003 the Fed's post-meeting statement primarily discussed the economic outlook and risks to that outlook. In the context of the illustrative model of Section 1.3.1, this can be formalized as clearer signals about  $\eta_t$ . Recalling the assumption that forecast revisions are larger for variables with clearer signals, this suggests a relatively high value of  $\hat{\sigma}_\eta$ —i.e., information effects are relatively larger (than monetary shocks) in the first regime. Therefore, the model predicts that forecast revisions to both variables will primarily reflect information shocks, which is borne out by the positive correlation in the top-left panel of Figure 1.5. Similarly, a relatively clearer signal about future interest rates implies a larger role for perceived future monetary shocks. This role becomes larger after 2003, as shown by the slightly negative correlation.

The change in correlation thus suggests that the variance of the underlying perceived shocks changed in 2003. One assumption allows this observation to be used to identify the model's structural parameters. Namely, by assuming that the slopes in the forecasting model ( $\phi$  and  $\gamma$ ) remain unchanged across the policy regions, Rigobon (2003) shows that only a unique pair of slopes can generate the data in both regimes (in a maximum-likelihood sense). I display that unique set of slopes in panel B of Figure 1.5 (delaying a discussion of their estimation to the next section). To see that these slopes are those most likely to generate the data, consider tilting the solid blue line down in the bottom-left panel so that it better fits the points in the first regime. This would necessarily cause it to fit the points in the top of the bottom-right graph.

Intuitively, then, the first sample primarily contributes to the estimation of the information shock, which, from the perspective of macroeconomic theory, should induce a positive correlation between these variables (which can be seen by shifting the red dashed line, the IS equation, and observing that all points are assumed to lie at the intersection of the two lines). Having learned about the information shock, the low to negative correlation in the second regime allows the monetary shock to be identified, since that regime exhibits variation in the two variables that is orthogonal to the information shock. In the extreme case in which the first regime contains only information shocks and the latter only monetary shocks, this intuition would be exact and the green lines

in panel A would provide structural slope coefficients. However, even with nonzero variances of each shock, the system can still be identified as long as the two shocks behave differently enough across the two regimes.

### 1.3.3 Formal Identification Assumptions

**Model** I estimate the following model:

$$\hat{y}_{\tau,t} = m_{11}\xi_{1,t} + m_{12}\xi_{2,t} \quad \xi_{1,t} \sim N(0, \sigma_{1,t}^2) \quad (1.9a)$$

$$\hat{i}_{\tau,t} = m_{21}\xi_{1,t} + m_{22}\xi_{2,t} \quad \xi_{2,t} \sim N(0, \sigma_{2,t}^2), \quad (1.9b)$$

where  $i_{\tau,t}$  is the 30-minute change in four-quarter ahead Eurodollar futures (section 1.2.4) and  $y_{\tau,t}$  is the high-frequency text-based proxy for output expectations (section 1.2.1). These observable variables are posited to be linear combinations of two independent Gaussian shocks,  $\xi_{1,t}$  and  $\xi_{2,t}$ . I assume that the shocks exhibit heteroskedasticity of the form

$$\sigma_{1,t} = \begin{cases} 1 & t \in \mathcal{R}_1 \\ \sigma_1 & t \in \mathcal{R}_2 \end{cases} \quad \sigma_{2,t} = \begin{cases} 1 & t \in \mathcal{R}_1 \\ \sigma_2 & t \in \mathcal{R}_2, \end{cases} \quad (1.10)$$

where, as described in section 1.3.2, the regimes,  $\mathcal{R}_1$  and  $\mathcal{R}_2$ , consist of all regularly-scheduled FOMC meetings running from May 1999 through June 2003, and August 2003 through December 2006, respectively. The normalization of the variances of the shocks is without loss of generality—their levels are not identified, only their relative levels between the two regimes. The model can be generalized to  $n$  variables and shocks, so for later reference I express the model in matrix form:

$$\underbrace{\hat{\mathbf{x}}_{\tau,t}}_{n \times 1} = \underbrace{\mathbf{M}}_{n \times n} \underbrace{\boldsymbol{\xi}_t}_{n \times 1} \quad \boldsymbol{\xi}_{\tau,t} \sim N(\mathbf{0}, \underbrace{\boldsymbol{\Sigma}_t}_{n \times n}). \quad (1.11)$$

where

$$\Sigma_t = \begin{cases} \Sigma_{(1)} & t \in \mathcal{R}_1 \\ \Sigma_{(2)} & t \in \mathcal{R}_2. \end{cases}$$

with the normalization  $\Sigma_{(1)} = \mathbf{I}_n$ .

The formal model just described is similar to the illustrative model used to generate intuition—this can be seen by comparing equations (1.9a) and (1.9b) to equations (1.8b) and (1.8a). Here, I posit that changes in expectations about observable variables are driven by two independent structural shocks. In the illustrative model, the shocks were given names and the coefficients were signed based on macroeconomic theory. Here, the coefficients are unrestricted and the shocks will be based on the historical episode that I use to estimate them.

**Assumption 1: Heteroskedasticity** The assumption of heteroskedasticity is crucial. Without it, the model would not be identified. Only three empirical moments would be available from the two Gaussian observable variables— $\text{var}(\hat{y}_{\tau,t})$ ,  $\text{var}(\hat{i}_{\tau,t})$ , and  $\text{cov}(\hat{y}_{\tau,t}, \hat{i}_{\tau,t})$ —but there would be four parameters to estimate,  $m_{11}, m_{12}, m_{21}$ , and  $m_{22}$ .<sup>28</sup> Rigobon (2003) developed a solution to this identification problem. The assumption of heteroskedasticity allows me to estimate those three empirical moments in both regimes, for a total of six moments. The model has six parameters to estimate:  $m_{11}, m_{12}, m_{21}, m_{22}, \sigma_1$ , and  $\sigma_2$ . Therefore, the system is just-identified, with moment

---

<sup>28</sup>Without loss, the variances of the shocks could be normalized to unity.

conditions given by

$$\text{var}_1 (\hat{y}_{\tau,t}) = m_{11}^2 + m_{12}^2 \quad (1.12a)$$

$$\text{var}_1 (\hat{i}_{\tau,t}) = m_{21}^2 + m_{22}^2 \quad (1.12b)$$

$$\text{cov}_1 (\hat{y}_{\tau,t}, \hat{i}_{\tau,t}) = m_{11}m_{21} + m_{21}m_{22} \quad (1.12c)$$

$$\text{var}_2 (\hat{y}_{\tau,t}) = m_{11}^2\sigma_1^2 + m_{12}^2\sigma_2^2 \quad (1.12d)$$

$$\text{var}_2 (\hat{i}_{\tau,t}) = m_{21}^2\sigma_1^2 + m_{22}^2\sigma_2^2 \quad (1.12e)$$

$$\text{cov}_2 (\hat{y}_{\tau,t}, \hat{i}_{\tau,t}) = m_{11}m_{21}\sigma_1^2 + m_{21}m_{22}\sigma_2^2, \quad (1.12f)$$

where the notation  $\text{var}_1 (\hat{y}_{\tau,t})$  is the variance of  $y_{\tau,t}$  in regime 1,  $\text{cov}_2 (\hat{y}_{\tau,t}, \hat{i}_{\tau,t})$  is the covariance of the observable variables in regime 2, and so on. Those moments—on the left-hand-side of each equation—can be estimated directly in the data. These six equations therefore have six unknowns—the parameters on the right-hand-side of each equation. The solution can be expressed analytically (see Rigobon (2003)). As solutions to quadratic equations, however, the associated expressions are not particularly enlightening.<sup>29</sup> An important condition for the system to be identified is that

$$\sigma_1 \neq \sigma_2. \quad (1.13)$$

In words, this rank condition states that the relative variances of each shock (between regimes) must be different. Suppose instead that the relative variances were equal, defining  $c \equiv \sigma_1 = \sigma_2$ . In this case, empirically, the six equations 1.12 would reduce to only three linearly independent equations, with (1.12a), (1.12b), and (1.12c) just  $c$ -scaled versions of (1.12d), (1.12e), and (1.12f). That parameter,  $c$ , could be (over-) identified, but not the  $m_{ij}$ .

**Assumption 2: Stable Effects** The second crucial assumption I make is that the  $m_{ij}$  coefficients remain unchanged across the two regimes. Without this assumption, four additional parameters (10

---

<sup>29</sup>In the (over-identified) case that one of the shocks is assumed to maintain the same variance between both regimes (Wright, 2012; Arai, 2017; Nakamura and Steinsson, 2018; Hébert and Schreger, 2017), an intuitive expression the “slope” coefficients  $m_{ij}$  emerges as essentially a change in OLS coefficients between the two regimes.

parameters total) would need to be estimated with only six equations. Empirically, this assumption can be tested by adding an additional regime and running over-identification ( $J$ ) tests—I present results of such a test alongside my estimates. To assess its theoretical validity, I consider the assumption within the context of linear rational (RE) expectations models. In full-information RE models there is no role for information effects, so I focus my discussion on linear RE models with imperfectly-informed agents.

I consider first the class of models discussed by Blanchard et al. (2013). In these models, all agents are imperfectly informed about the economy’s state variables. Models of this class take the form

$$\mathbf{A} \mathbf{x}_t + \mathbf{B} \mathbb{E}_t[\mathbf{x}_{t+1}] + \mathbf{C} \boldsymbol{\xi}_{t|t} = \mathbf{0} \quad (1.14)$$

where  $\mathbf{x}_t$  is a vector of observable macroeconomic variables,  $z_{t|\tau}$  denotes the mathematical expectation of  $z_t$  given information at time  $\tau$ , and  $\boldsymbol{\xi}_t$  is a vector of mutually independent structural shocks that evolves according to  $\boldsymbol{\xi}_t = \mathbf{D}\boldsymbol{\xi}_{t-1} + \mathbf{F}\boldsymbol{\zeta}_t$ , where  $\boldsymbol{\zeta}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_\zeta)$ .<sup>30</sup> At the beginning of each period  $t$ , but before making decisions about  $\mathbf{x}_t$ , all agents receive the same noisy signal of the structural shock of the form  $\mathbf{s}_t = \mathbf{G}\boldsymbol{\xi}_t + \mathbf{H}\boldsymbol{\nu}_t$ , where  $\boldsymbol{\nu}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_\nu)$ . Agents use this signal to form expectations about  $\boldsymbol{\xi}_t$  using the Kalman filter (i.e., agents have RE). Agents are thus imperfectly informed in a symmetric way—they all share the same information set when making decisions.

In appendix C, I show that the model in equation (1.14) admits a solution of the form

$$\mathbf{x}_t = \mathbf{J}\boldsymbol{\xi}_{t|t} \quad (1.15)$$

along with a law of motion for perceived shocks,  $\boldsymbol{\xi}_{t|t}$ , where  $\mathbf{J}$  depends on neither  $\boldsymbol{\Sigma}_\zeta$  nor  $\boldsymbol{\Sigma}_\nu$ . Equations 1.15 and the law of motion for  $\boldsymbol{\xi}_t$  reveal the result: the mapping between forecast revisions of

---

<sup>30</sup>The assumption that the shocks follow a VAR(1) is not restrictive—any finite VARMA can be re-written as a VAR(1). Adding the vector of lagged endogenous variables to equation (1.14) (i.e.,  $\mathbf{x}_{t-1}$ ) slightly changes the language, but not results, of this discussion. Specifically, rather than forecast revisions of endogenous variables being related linearly to revisions of structural shocks, it’s forecast revisions of *surprises* (i.e., reduced form residuals) of endogenous variables that are related linearly to revisions of structural shocks.

observable variables and forecast revisions of structural shocks is fixed and linear:

$$\mathbb{E}_t[\mathbf{x}_{t+k}] - \mathbb{E}_{t-1}[\mathbf{x}_{t+k}] = \mathbf{J}\mathbf{D}^k[\boldsymbol{\xi}_{t|t} - \boldsymbol{\xi}_{t|t-1}].$$

This relationship resembles the relationship between reduced-form and structural shocks in the parlance of structural VARs, another macroeconomic model in which the assumption of constant  $m_{ij}$  is valid.

The key ingredient for a linear RE model to feature a constant mapping between observable variables and structural shocks is that the model’s structural equations can be solved independently of agents’ Kalman filtering problem. The result would continue to hold, therefore, in models in which agents also observe the endogenous variables,  $\mathbf{x}_t$ . An important case in which this independence breaks is in linear RE models with dispersed information (as in the “islands” model of Lucas (1972)). Within the context of these models, it then becomes an empirical question: To what extent does the dependence of  $m_{ij}$  on the variance of the structural shocks affect the ability of the shocks to be recovered using the heteroskedasticity-based assumptions? In appendix J.2.2, I ask this question within the context of the model of Fed “signalling effects” posited and estimated by Melosi (2017), which features agents with dispersed information and realistic modeling of Fed signaling. Specifically, I simulate the “introduction of forward guidance” experiment within the model. I find that the shocks identified by heteroskedasticity uncover their structural counterparts remarkably well. In appendix J.2.3, I also discuss why the assumption of dispersed, rather than common, knowledge eliminates the aforementioned independence in the context of a simple asset pricing model that follows Townsend (1983).

**Assumption 3: Naming** In the illustrative model, the two structural shocks,  $\hat{\eta}_{\tau,t}$  and  $\varepsilon_{\tau,t}$ , were given names consistent with macroeconomic theory. The model of equations (1.9a) and (1.9b) instead is driven by two structural shocks,  $\xi_{1,t}$  and  $\xi_{2,t}$ , with no natural names. Members of the high-frequency identification literature typically impose identification restrictions on the  $m_{ij}$  coef-

ficients.<sup>31</sup> In this paper, I take the view that the  $m_{ij}$  are exactly the objects that require unrestricted estimation, since they ultimately determine the nature of the estimated shocks. Instead, I place assumptions on the size of the underlying shocks, motivated by the historical evidence. Here, the illustrative model and the historical evidence provide a natural solution: The  $\xi$  shock that I call the “information shock” is the shock whose variance becomes relatively smaller in the second regime. Formally, recalling that  $\sigma_i = \text{var}_2(\xi_{i,t})/\text{var}_1(\xi_{i,t})$ , I name the information shock as

$$\hat{\xi}_t^I \equiv \text{information shock}_t \equiv \arg \min_{\xi_{i,t}}[\sigma_i]. \quad (1.16a)$$

Conversely,

$$\hat{\xi}_t^M \equiv \text{monetary path shock}_t \equiv \arg \max_{\xi_{i,t}}[\sigma_i]. \quad (1.16b)$$

I therefore name my shocks based on the size of their relative variance across the regimes. It then becomes an empirical question—not an assumption—whether these shocks have theoretically-consistent effects on high-frequency forecast revisions. This can be determined through the estimates of the  $m_{ij}$ . I carry out the estimation in the next section.

#### 1.3.4 Estimation

I estimate the parameters using GMM and calculate bootstrapped standard errors and confidence intervals. In appendix G I describe the bootstrap procedure in detail. The procedure is mostly standard—nearly identical to that of Hébert and Schreger (2017)—though I design a method to handle the fact the columns of  $\mathbf{M}$  are only identified up to order and sign.

I present the estimates in table 1.3. Panel A contains the estimated structural impact matrix,  $\mathbf{M}$ , and panel B shows the estimates of  $\Omega_2\Omega_1^{-1}$ —the variance of each of the structural shocks in  $\mathcal{R}_2$  relative to  $\mathcal{R}_1$ .<sup>32</sup> Following the discussion of equation (1.16), I name the “information shock” based

---

<sup>31</sup>For example, Kuttner (2001) and Gürkaynak et al. (2005) impose zero restrictions on the  $m_{ij}$ , while (Jarociński and Karadi, 2020) impose sign restrictions.

<sup>32</sup>Recall that I am allowed the normalization of two parameters without loss—here I have normalized  $\Omega_1$  to be a  $2 \times 2$  identity matrix.

Table 1.3: Heteroskedasticity-Based Estimates

Panel A: Structural Impact Matrix,  $\mathbf{M}$

	<i>Response of High-Frequency Forecast Revisions</i>	
	Eurodollars (1Y)	GDP
Monetary Path Shock	-0.39 [-1.17, 0.13]	0.78 [0.55, 1.28]
Information Shock	0.99 [0.71, 1.43]	0.56 [0.02, 1.29]

Panel B: Relative Variance in Regime 2,  $\Sigma_2 \Sigma_1^{-1}$

Monetary Path Shock	Information Shock
1.37 [0.69, 2.00]	0.63 [0.38, 1.58]

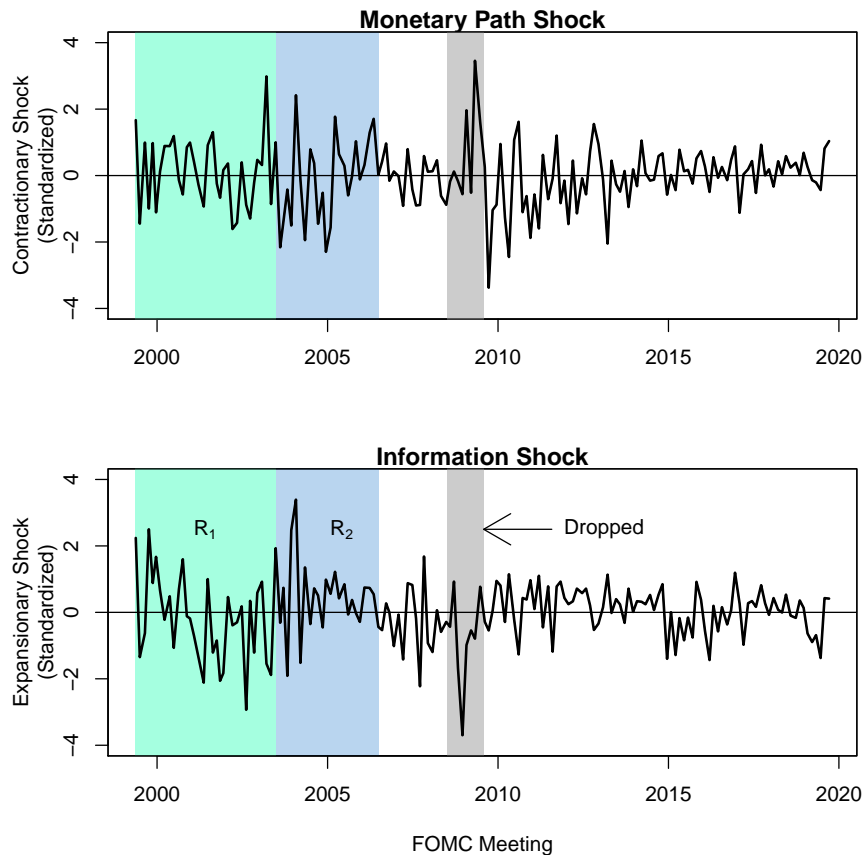
This table shows the estimates of the system in equation (1.12). Panel A shows the structural impact matrix, which shows the effect that each structural shock (in the rows) has on forecast revisions made around FOMC announcements (in the columns). The shocks are normalized to have unit variance in the first regime (2000–2003). Panel B shows the variance of each shock in regime 2 relative to its variance in regime 1. 90% equal-tailed (studentized) confidence intervals are produced using 999 bootstrap replications, where bootstrap samples are stratified by regime.

on the fact that it is relatively smaller in the second regime than in the second. This is consistent with the notion that the Fed’s statement in the first regime focused exclusively on describing the state of the economy. Because I use interest rate expectations at a one-year horizon, I call the second shock a “monetary path shock”—an exogenous shock to the expected path of interest rates. Again, once the Fed starts to communicate explicitly about future interest rates, markets are able to make more-informed (larger) updates about expected future monetary shocks. Formally, I can reject the null hypotheses that the variance of the monetary shock is relatively smaller than the variance of the demand shock in the second regime with 90 percent confidence.<sup>33</sup>

Having named the shocks based on their relative variances, the estimates of  $\mathbf{M}$  serve as the first

<sup>33</sup>Let  $\delta$  be the difference between the variance of the monetary and information shock. The 90% confidence interval for this test is  $(x, \infty)$ , where  $x$  is  $\hat{\delta} - \text{s.e.}(\hat{\delta})G_n^{*-1}(0.9)$ , where  $\hat{\delta}$  is the point estimate of  $\delta$ ,  $\text{s.e.}(\hat{\delta})$  is the standard deviation across bootstrap replications, and  $G_n^*$  is the bootstrap distribution of t-statistics  $(\hat{\delta}^* - \hat{\delta})/\text{s.e.}(\hat{\theta}^*)$ . Note that estimation  $\hat{\theta}^*$ , requires an inner bootstrap (bootstrapping the bootstrap) for which I also use 999 replications. Empirically,  $x = 0.002$ .

Figure 1.6: Time Series of Estimated Monetary and Information Shocks



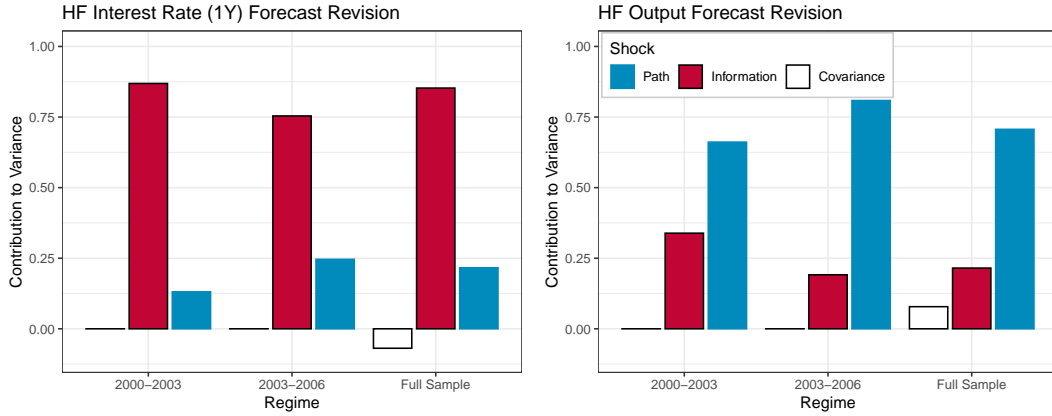
The top panel shows the identified monetary policy path shock described in section 1.3.4, and the bottom panel shows the identified information shock. Both shocks are standardized to have unit standard deviation over the plotted sample. The first two shaded areas present the regimes used for identification, described in section 1.3.2. The third, gray, shaded area highlights observations that are dropped for most analysis, discussed in section 1.2.4.

check on the shocks. A monetary policy shock that lowers interest rates by 0.37 standard deviations *increases* GDP expectations by 0.76 standard deviations. Conversely, a monetary policy shock that lowers rates by 1 standard deviation *decreases* GDP expectations by 0.57 standard deviations. These impact responses are consistent standard New-Keynesian macroeconomic theory and the illustrative model. The time series of the estimated shocks, given by  $\mathbf{M}^{-1}\hat{\mathbf{x}}_t$ , are shown in figure 1.6.<sup>34</sup>

To have a better sense of how the identified monetary and information shocks relate to inter-

<sup>34</sup>Here I make a less demanding form of an “invertibility” assumption, as discussed in the structural-VAR literature (see, e.g., Chahrour and Jurado (2021) or Fernández-Villaverde et al. (2007)). The illustrative model and discussion surrounding equation (1.15) provided the justification for its validity here. Note that my context—recovering *perceived shocks* from expectations data—is different from the VAR context, in which one seeks to identify structural shocks from reduced-form VAR residuals.

Figure 1.7: Variance Decomposition of HF Forecast Revisions



This figure shows which portion of the variance of HF interest rate forecast revisions,  $\hat{i}_t$  (left panel), and HF output forecast revisions,  $\hat{y}_t$  (right panel), can be explained by the monetary  $\hat{\varepsilon}_t$  and information  $\hat{\eta}_t$  shocks. Recalling from equations (1.9a) and (1.9b) that  $\hat{i}_t = \phi_{i\varepsilon} \hat{\varepsilon}_t + \phi_{i\eta} \hat{\eta}_t$  and  $\hat{y}_t = \phi_{y\varepsilon} \hat{\varepsilon}_t + \phi_{y\eta} \hat{\eta}_t$ , the red outlined boxes for  $j \in \{i, y\}$  show  $\phi_{j\eta}^2 \text{var}(\hat{\eta}_t) / \text{var}(\hat{j}_t)$ ; the blue non-outlined boxes show  $\phi_{j\varepsilon}^2 \text{var}(\hat{\varepsilon}_t) / \text{var}(\hat{j}_t)$ ; and the white-outlined boxes show the remaining covariance term (which is not restricted to be null in the full sample).

est rate and GDP forecast revisions, figure 1.7 shows the variance decomposition of the forecast revisions. The calculation during the period 1999–2006 follows immediately from the estimated system—the variance of forecast revisions is the weighted (by  $\mathbf{M}$ ) sum of the variances of the underlying shocks. Both during the estimation period and in the full sample, it is apparent that the bulk of interest rate forecast revisions are driven by the information shock. This is not a surprising finding given the puzzling estimated effects of interest rate surprises, but these estimates allow me to quantify the extent of the contamination problem. Roughly about 70% of interest rate surprises (recall, these are traditional HF estimates of monetary policy shocks) are made up of information. This contamination was less severe during 2003–2006 than it was in 2000–2003, consistent with the illustrative model.<sup>35</sup>


<sup>35</sup>To see this, notice that the variance of interest rate surprises is:

$$\text{var}(\hat{i}_t) = \underbrace{\left(\frac{\phi}{1 + \gamma\phi}\right)^2}_{\text{Var. due to info.}} \nu_\eta + \underbrace{\left(\frac{1}{1 + \gamma\phi}\right)^2}_{\text{Var. due to MP}} \nu_\varepsilon.$$


The relative contribution of information effects to interest rate forecast revisions is therefore  $\phi^2 \left(\frac{\nu_\eta}{\nu_\varepsilon}\right)$ . In the model, a regime in which post-meeting statements focus on the Fed’s outlook for the state of the economy is akin to increasing the precision of the information signal, i.e., lower  $\sigma_{n,\eta}$ , thus higher  $\nu_\eta$ , which increases the relative contribution of information effects to interest rate surprises. Thus, a prediction of the model is that relative to 2003–2006, interest rate surprises during 2000–2003 should be more contaminated by information effects.

Figure 1.8: Baseline Macroeconomic Effects

Panel A: Responses to a Contractionary Monetary Path Shock ( $\beta_k^M$ )



Panel B: Responses to an Expansionary Information Shock ( $\beta_k^I$ )



This figure plots the estimates of equation (1.17). Moving from left to right, the left-hand-side variables are  $k$ -period differences in 100 times the log of industrial production (FRED mnemonic INDPRO), 100 times the log of the core PCE prices index (PCEPILFE), and the shadow Federal Funds rate of Wu and Xia (2016) (in percent). The right-hand-side contains 12 lags of these variables, as well as the estimated monetary path shock and information shock, standardized to have unit standard deviation over the full sample period, which runs from May 1999 to October 2019, excluding July 2008–July 2009. Both shocks are also scaled to increase the 1-day change in the one-year treasury yield on impact—thus the monetary shock is contractionary, and the information shock is expansionary. I exclude months with no FOMC meeting. Confidence intervals are calculated using heteroskedasticity and autocorrelation-consistent asymptotic standard errors with the automatic lag selection method of Newey and West (1994), as implemented by Zeileis et al. (2020) and Zeileis (2004).

The “full sample” column in each panel serves as a check on the assumption that the shocks are uncorrelated. During the 1999–2006 period the shocks are forced to be orthogonal. During the remainder of the sample, this orthogonality condition is not imposed, so a covariance term enters the expression for the variance of forecast revisions. The covariance between the shocks contributes negligibly to the variance of observed forecast revisions, suggesting that the orthogonality restriction was warranted.

## 1.4 Effects of Monetary Policy and Information

I now turn to estimating the effects of monetary policy and Fed information. In Section 1.4.1, I present my baseline estimated effects on macroeconomic outcomes and expectations. My specifications are purposefully simple. I use local projections (Jordà, 2005) to estimate macroeconomic effects and OLS to estimate the response of macroeconomic expectations. Since part of my objective in this paper is to provide a credible and portable set of shocks, it is important that my shocks have sensible effects without depending on particular controls. In Section 1.4.2, I compare my estimated effects with those found using other measures in the literature. Along these lines, in section 1.4.3 I put my estimated effects in policy-relevant terms by estimating instrumental-variables versions of my baseline specification.

### 1.4.1 Baseline Results

**Effects on Macroeconomic Outcomes** In order to estimate the effects of my shocks on macroeconomic outcomes, I estimate local projections of the form

$$y_{t+k} - y_{t-1} = \alpha_k + \beta_k^M \xi_t^M + \beta_k^I \xi_t^I + \sum_{\ell=1}^L \Gamma'_{\ell,k} \Delta \mathbf{y}_{t-\ell} + \xi_{k,t}, \quad (1.17)$$


where  $\xi_t^M$  and  $\xi_t^I$  are the estimated monetary and information shocks.<sup>36</sup> The left-hand-side variable,  $y_t$ , is either the log of the industrial production index, the log of the core PCE price index, or the shadow Federal Funds rate of Wu and Xia (2016). The vector  $\mathbf{y}_t$  contains the three  $y_t$  variables, of which I include  $m = 12$  lags. I estimate the equation at monthly frequency, using the sample described in Section 1.2.4. The shocks are normalized to have unit standard deviation over the regression sample, with a positive effect on interest rates (the high-frequency change in the 4-quarter Eurodollar future). Thus, the coefficients of interest are  $\beta_k^M$ —the effects of a contractionary monetary policy shock—and  $\beta_k^I$ , the effects of an expansionary information shock.

---

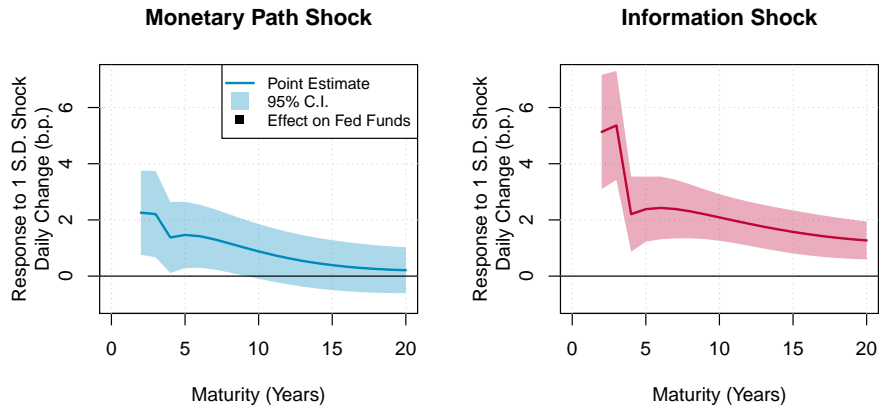
<sup>36</sup>Note, as described in Section 1.3.4, that the shocks are essentially uncorrelated over the sample, so their joint inclusion does not affect the point estimates, only the precision.

Figure 1.9: One-Day Changes in Yields

Panel A: Nominal Yields



Panel B: Real Yields



This graphs show the estimated slope coefficient from univariate regressions of 1-day changes in nominal (Panel A) and real (Panel B) interest rates of maturity  $m$  ( $\Delta i_t^m$ ) on the monetary path shock (left panel) and the information shock (right panel), using my baseline sample. Regarding nominal rates, for maturities of at least 1 year, I use data from Gürkaynak et al. (2007) The Federal Funds, three- and six-month yields are from FRED, with mnemonics EFFF, DGS3MO, and DGS6MO, respectively. The zero lower bound leads to enormous standard errors in the Fed Funds rate regression, so I omit them. I take real rates from Gürkaynak et al. (2010). Confidence intervals are computed using robust standard errors.

Figure 1.8 contains the estimated coefficients. The top panel shows the effect of a contractionary monetary policy shock. Encouragingly, both industrial production and inflation decrease. The delayed peak responses of these variables—at about 36 months—is consistent with the notion that monetary policy works with “long and variable lags.”

The response of the shadow rate to a monetary policy shock is less clear-cut. In high frequency, a monetary shock leads to an upward revision of interest-rate expectations. As shown in Panel A

Figure 1.9, nominal yields several years into the term structure also increase in a 1-day window around the policy announcement in response to a monetary shock. The effects at a monthly frequency are imprecisely estimated and fluctuate around zero. Figure 1.10 investigates whether this finding is robust across alternative measures of nominal rates. Specifically, I re-estimate equation (1.17) using longer-term interest rates in place of the shadow rate. While the estimates continue to be imprecise, the negative effect on the interest rate diminishes as increasingly longer-term interest rates are used.

Panel B of Figure 1.8 shows the effects of an expansionary information shock. While the responses are more immediate than responses to a monetary shock, the peak effects on industrial production and prices are similar in magnitude (in absolute value) to the peak effects of a monetary shock. This suggests that the Fed plays an important role not only in setting interest rates, but also in the provision of macroeconomic information.<sup>37</sup>

The relative responses of industrial production and inflation also give a hint as to the nature of the shocks about which information is revealed. The positive comovement in response to an information shock suggests that the information primarily concerns demand-type factors. In Section 1.5.2 I provide further high-frequency evidence in support of this conjecture.


In Appendix H, I examine the robustness of the estimates of equation (1.17). I first add months with no FOMC meetings, setting the shocks in those months to zero. I next stop the estimation in 2015 in order to (i) obtain a consistent sample across all horizons and (ii) drop COVID observations. Similarly, I include a specification that retains Great Recession observations. I test the functional form of the controls by first controlling for no lags, then for 24 lags, and then add a linear trend. Finally, I use a 10-word word list  $\mathcal{Y}$  for GDP to construct the shocks. Across these specifications, the responses are largely similar to the estimates shown in Figure 1.8.

**Effects on Macroeconomic Expectations** I close this section by studying the effects of monetary and information shocks on macroeconomic expectations. Specifically, I regress 1-month

---

<sup>37</sup>These estimates should prove useful for informing models of Fed information effects, e.g., that of Lepetyuk et al. (2021).

Figure 1.10: Effects of Monetary Shocks: Using Longer-term Interest Rates



These graphs show the response of industrial production, core PCE inflation, and interest rates to the identified monetary shock. These are estimates of equation (1.17), except that I replace the shadow rate of Wu and Xia (2016) with either the two-year treasury yield (solid line, FRED mnemonic DGS2), the five-year treasury yield (dashed line, DGS5), and the twenty-year treasury yield (dotted line, DGS20). See the note to table 1.8 for details regarding units, sample period, and confidence intervals.

changes in macroeconomic expectations from the Blue Chip survey on the two shocks. The regression takes the form

$$\mathbb{X}_{t|t+1} - \mathbb{X}_{t|t} = \alpha + \beta^M \xi_t^M + \beta^I \xi_t^I + e_t, \quad (1.18)$$

where  $\mathbb{X}_{t|t}$  is constructed analogous to equation (1.3), except now I study the responses of real GDP growth (as above), CPI inflation, and the unemployment rate. I present the results in Table 1.4. Again, the shocks are scaled to increase interest rates in high frequency—thus, the first row represents the effects of a contractionary monetary policy shock and the second represents the effects of an expansionary information shock.

The signs of the responses of all variables are consistent with both the macroeconomic effects estimated above, and standard macroeconomic theory. The information shock has larger (and more precisely estimated) effects on real variables, which highlights another way in which the Fed plays an important role in the provision of macroeconomic information. The information shock also induces a positive correlation between these variables and inflation, which again suggests that this information concerns demand-type factors.

#### 1.4.2 Comparison with Previous Estimates

While my paper is not the first that seeks to provide a credible measure of monetary policy shocks, my shocks differ in that they robustly exhibit theoretically consistent macroeconomic ef-

Table 1.4: Shocks and Macroeconomic Expectations

	GDP Growth	CPI Inflation	Unemp. Rate
Path Shock	-0.0126 (-1.36)	-0.0152 (-2.53)	0.0206 (2.14)
Info. Shock	0.0389 (3.99)	0.0138 (1.43)	-0.0256 (-3.69)
Observations	131	131	131
$R^2$	0.116	0.0614	0.105
$ \overline{\text{LHS}} $	0.0775	0.0532	0.0625

This table shows estimates of equation (1.18). The left-hand-side is 1-month forecast revision of the variables listed atop the columns over the next three forecasting horizons, presented in equation (1.3), from the Blue Chip survey in percentage points. The estimated monetary path shock and information shock are standardized to have unit standard deviation over the full sample period. The sample consists of all regularly scheduled FOMC meetings between May 1999 and October 2019, excluding July 2008–July 2009, that occur after the first week of the month. The row  $|\overline{\text{LHS}}|$  is the average absolute Blue Chip revision in the relevant column over the regression sample.

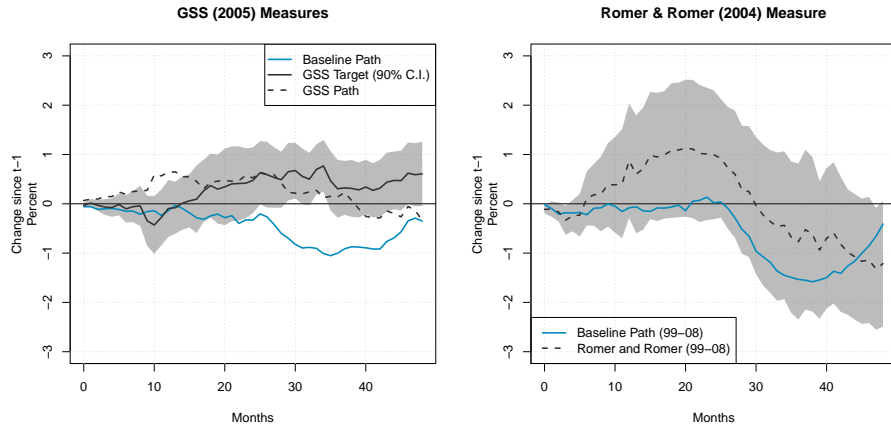
fects. In this section, I compare the estimated effects of alternative measures of monetary shocks found in the literature. I estimate the effects using the specification in equation (1.17), in which I replace  $(\beta_k^M \xi_t^M + \beta_k^I \xi_t^I)$  with a single measure of monetary policy shocks.

### Traditional Measures

I first study the estimated effects of monetary policy using shocks that are pervasive in empirical macroeconomics. Specifically, I study the high-frequency shocks of Gürkaynak et al. (2005).<sup>38</sup> Those authors decompose high-frequency changes in interest rate expectations for maturities out to 1 year into a “target” and “path” factor. The shocks are identified and named with the assumption that the path factor has no effect on the current-meeting Federal Funds rate surprise. I updated these measures using the tick data described in Section 1.2.4. I also study the shocks of Romer and Romer (2004). Those shocks are constructed as the component of the change in the Federal Funds rate that cannot be predicted from the Fed staff’s (“Greenbook”) forecasts. I updated this series through 2012 (though I stop my estimation at the start of the 2008–2015 zero lower bound period).

<sup>38</sup>Other well-known high-frequency shocks are those of Kuttner (2001), Nakamura and Steinsson (2018), and Gertler and Karadi (2015). Those all differ by the maturity of the underlying interest rates. The shocks of Gürkaynak et al. span the maturity spectrum used by those other papers.

Figure 1.11: Response of Industrial Production to Traditional Contractionary MP Shocks



These graphs show the response of industrial production to various measures of monetary shocks. The blue line in the left panel is identical to that in Figure 1.8. All other lines are computed by estimating equation (1.17), replacing  $(\beta_k^M \xi_t^M + \beta_k^I \xi_t^I)$  with  $\beta \text{shock}_t$ , where  $\text{shock}_t$  is either the target or path factor of Gürkaynak et al. (2005) (left panel), or the shock of Romer and Romer (2004) (right panel), described in the text. In the right panel I re-estimate the effects of my shocks with an abbreviated sample (1999–2008) for comparison with the effects of the shock of Romer and Romer. All shocks are standardized to have unit standard deviation over the sample period, and increase the 1-day change in the 1-year treasury on impact—they are all thus 1-standard deviation contractionary shocks. See the note to Table 1.8 for details regarding units, sample period, and standard errors.

I focus on the effects on industrial production and present the results in Figure 1.11. All shocks are scaled to increase the 1-day change in the 1-year treasury yield, so the results can be interpreted as the responses to estimates of contractionary monetary policy shocks. The left panel shows the type of puzzling evidence suggested in the introduction: A contractionary high-frequency shock leads to increases in industrial production. This is the case whether regardless of the horizon of the interest-rate surprise (i.e., the target or path factor).

In the right panel of figure 1.11, I show the responses using the measure constructed by Romer and Romer (2004). I also re-estimate the effects of my monetary shock using the same (1999–2008) sample. Romer and Romer’s responses are fairly noisy and fluctuate around zero, though all statistically significant estimates have the theoretically consistent sign. This last finding suggests that somehow adjusting for the Fed’s private information is useful for estimating theoretically-consistent effects of monetary policy.

## New Measures of Monetary Shocks

The findings in the previous study also inspired other authors to estimate “information free” measures of monetary policy shocks. The predominant method of point-identifying information free measures was developed by Miranda-Agrippino and Ricco (2021), who orthogonalize interest rate surprises to the Fed staff’s forecasts—essentially a combination of the approach of Kuttner (2001) and Romer and Romer (2004).<sup>39</sup> I downloaded the shocks of Miranda-Agrippino and Ricco (2021) directly from the first author’s website.

The left panel of figure 1.12 contains the result. The estimated effects of Miranda-Agrippino and Ricco (2021) exhibit theoretically consistent effects, with much-more precise estimates than those of Romer and Romer (2004). In fact, the responses to the Miranda-Agrippino and Ricco shock are similar to my baseline estimates. At a first glance both seemingly-plausible identification approaches produce similar results.

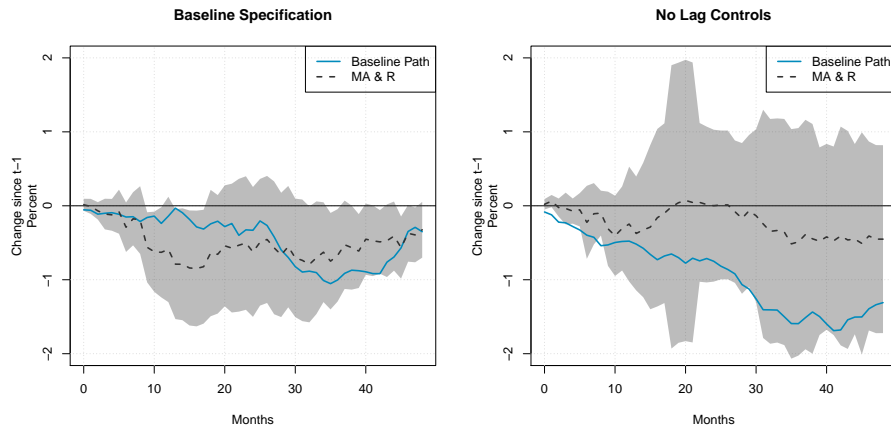
Important differences arise when the specification is substantially simplified. In the right panel of Figure 1.12, I remove any controls from the regression in equation (1.17). The estimated effects using Miranda-Agrippino and Ricco’s shock diminish substantially in magnitude and precision. The fact that my monetary shock, in contrast, provides similar estimates using this simpler specification speaks to its portability to other contexts. Put differently, my shocks can be used directly to estimate the effects of monetary policy, without requiring additional controls.<sup>40</sup>

---

<sup>39</sup>Hoesch et al. (2020) and Zhang (2020) take a similar approach.

<sup>40</sup>This result is unsurprising given the construction of Miranda-Agrippino and Ricco (2021)’s measure. It is important to note that information effects arise because the Fed and public have different information sets. Therefore, to “remove” information effects, it is imperative to control for the *difference* between the Fed’s and public’s information. By only controlling for the Fed’s information, Miranda-Agrippino and Ricco’s shock “leaves behind” the public’s information. Thus, (at best) their measure is contaminated by (classical) measurement error that attenuates their estimated effects. The lags of observable variables in the expanded local projection likely span the public’s information set, which thus allows for the unbiased estimates shown in the left panel.

Figure 1.12: Response of Industrial Production to “Information Free” Contractionary MP Shocks



These graphs show the response of industrial production to the shock of Miranda-Agrippino and Ricco (2021). The left panel includes 12 lags of observable variables as controls, while in the right panel no lags are included. See Figure 1.12 for other details.

### 1.4.3 Comparison with Previous Estimates: Magnitudes

The estimated shocks have no interpretable magnitudes because they are linear combinations of two variables with different units.<sup>41</sup> The estimates presented thus far, therefore, speak to average observed influence of the Fed on macroeconomic aggregates, but do not put those effects in policy-relevant terms. In this section I follow the work of Gertler and Karadi (2015) and Ramey and Zubairy (2018) and estimate instrumental variables versions of equation (1.17). This specification, which I refer to as “LP-IV,” takes the form

$$y_{t+k} - y_{t-1} = \alpha_k + \beta_k p_t + \sum_{\ell=1}^L \Gamma'_{\ell,k} \Delta \mathbf{y}_{t-\ell} + \xi_{k,t}, \quad (1.19)$$


where  $p_t$  is the policy-relevant variable of interest, for which I will instrument with the relevant policy shock. For estimating the effects of monetary policy, interest rates are a natural candidate, so I follow Gertler and Karadi (2015) and set  $p_t$  to the level of the 1-year Treasury yield. The first-stage  $F$ -statistic here is 14.<sup>42</sup> The policy-relevant variable or estimating the effects of information

<sup>41</sup>One of those variables—the GDP forecast revision proxy—no interpretable units in the first place.


<sup>42</sup>For the effects of monetary policy, I set  $L = 6$ . For the effects of information, I set  $L = 12$ . For reasons still under investigation, there are a handful of outlier observations that emerge when setting  $L \approx 12$  in the monetary regression. These outliers cause a few of the point estimates, and their standard errors, to become orders of magnitude larger than those shown here. With outliers removed, the results are nearly identical to those shown here.

Figure 1.13: LP-IV Estimates

Panel A: Effects of Monetary Policy (1-year Treasury Yield)



Panel B: Effects of Information Provision (GDP Forecast Revision)



This figure plots the estimates of  $\beta_k$  from equation (1.19), where  $p_t$  and its instrument are described in the text. Moving from left to right, the left-hand-side variables are  $k$ -period differences in 100 times the log of industrial production (FRED mnemonic INDPRO), the 100 times the log of the core PCE prices index (PCEPILFE), and the 1-year Treasury yield in period  $t + k$  in percent (DGS1, though the top-right panel shows the result of the level of 1-year Treasury yield, to make the effect on the instrument clear). Panel A has six lags of the first-differences in these three variables, panel B has twelve. The sample runs from May 1999 to October 2019, excluding July 2008–July 2009. I exclude months with no FOMC meeting. Confidence intervals are calculated using heteroskedasticity and autocorrelation-consistent asymptotic standard errors with the automatic lag selection method of Newey and West (1994), as implemented by Baum et al. (2010).

provision are less obvious, but following the framework of the rest of the paper, I posit GDP forecast revisions (from the Blue Chip) as a reasonable measure. Here the first-stage  $F$ -statistic is 24—an unsurprising result, given the results in Table 1.4, whose “GDP Growth” variable is the  $p_t$  used here.

Figure 1.13 contains the results. To use the terminology of Coibion (2012)—who compares and reconciles the estimated effects of monetary policy from several prominent papers—the effects of monetary policy are big. The peak responses are similar in magnitude (about a 5% drop to a 1 p.p. increase in nominal rates) and timing (delayed by less than a year) as those estimated by Romer and Romer (2004). As Coibion points out, these estimates are on the upper end of estimates found in the literature. They are substantially larger than the effects found in the existing literature that uses high-frequency shocks in an LP-IV setting (Gertler and Karadi, 2015; Miranda-Agrippino

and Ricco, 2021). Figure 1.9 helps to make sense of these big effects. There, these increases in nominal interest rates are seen to translate nearly completely to increases in real interest rates.

While little existing empirical work serves as a reference point for understanding the effects of information provision, figure 1.13 suggests that the effects are big as well. A 1 p.p. increase in GDP expectations leads to a nearly 20% increase in industrial production. Care should be taken in interpreting these results, however. The average absolute GDP forecast revision is less than 0.1 percentage point, so these estimates are largely extrapolating outside of historical experience. The response of interest rates can help put the results in perspective—those increase by about 4p.p., about five-times less than industrial production.<sup>43</sup> In this sense, the results are also similar in magnitude to the effects of monetary policy.<sup>44</sup>

## 1.5 Expanding the Picture: Short-term Rates and Supply Factors

As discussed in the introduction, expanding the set of observable variables can give the econometrician greater insight into the nature of the shocks perceived by markets. In this section, I expand the set of variables along two dimensions. In Section 1.5.1, I study the role of short-term interest-rate surprises. The regimes I consider allow me to identify an additional shock—a monetary target shock—and study its behavior over time. In summary, following the intuition suggested by Ramey (2016), monetary policy shocks of this kind largely disappeared following the introduction of forward guidance. I also extend my model to include dynamic shocks in order to show how, simultaneously, path shocks can become larger while target shocks become smaller.

In Section 1.5.2, I present an estimate of high-frequency inflation forecast revisions. Around Fed announcements, inflation forecast revisions are unconditionally positively correlated with output forecast revisions. This is also true conditional on the estimated information and monetary path shocks. Both pieces of evidence suggest that markets do not learn about supply-type factors from

---

<sup>43</sup>I do not use nominal interest rates as my policy-relevant instrument for information effects for primarily for practical reasons—I have no first stage.

<sup>44</sup>These large estimated effects of the information shocks intuitively suggest that there is some “real” information content of these shocks—the canonical interpretation of information effects—and not a measure of macroeconomic disagreement, as in Sastry (2021).

the Fed. I also discuss why the regimes I use do not allow me to separately identify a “demand information” shock and a “supply information” shock. Intuitively, I can only identify shocks whose variance is induced to change across regimes; there is little evidence to support the notion that the introduction of forward guidance altered the relative precision of supply vs. demand signals.

### 1.5.1 Short- vs. Long-term Interest Rates

In their original incarnation, high-frequency monetary shocks were the surprise component of the Fed’s current-meeting interest-rate decision, in which the surprise was relative to expectations formed shortly before the decision was announced. This is the measure proposed by Kuttner (2001). Later work by Gürkaynak et al. (2005), cognizant of the 2003 introduction of forward guidance, sought to separately identify shocks to the interest-rate target from shocks to the interest-rate path. In this section I revisit this distinction, adding to my system the surprise in the current-meeting Fed Funds rate described in Section 1.2.4.

In Table 1.5, I show the estimates from estimating the moment conditions from equation (1.12), where now my observable variables are given by

$$\hat{\mathbf{x}}_t \equiv \left[ \hat{\text{FF}}_t \quad \hat{i}_t \quad \hat{y}_t \right]', \quad (1.20)$$

where  $\hat{\text{FF}}_t$  is the 30-minute change in current-month Federal Funds rate expectations (the Kuttner measure),  $\hat{i}_t$  is the 30-minute change in the 4-quarter-ahead Eurodollar futures price, and  $\hat{y}_t$  is the high-frequency text-based proxy for output expectations (Section 1.2.1). The regimes are unchanged.

Before discussing the identified “monetary target shock,” it is useful to examine the top two-by-two quadrant of panel A. There, notice that the relationship between the path shock and information shock with GDP and interest rate forecast revisions is nearly unchanged from Table 1.3.<sup>45</sup> This highlights the robustness of the identified information and path shocks. Neither shock has much of

---

<sup>45</sup>Recall that the sign of the rows of Table 1.5 are not identified (only the relative signs between elements of a row).

Table 1.5: Identifying a Target Shock

Panel A: Structural Impact Matrix,  $\mathbf{M}$

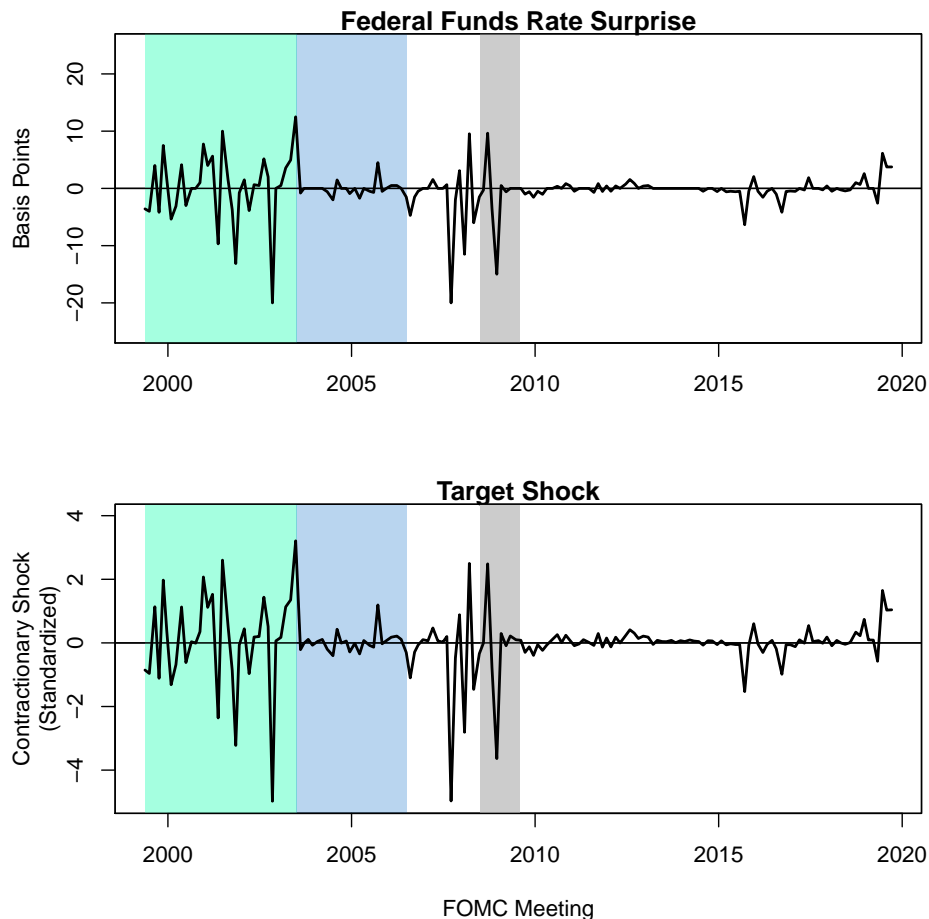
	<i>Response of High-frequency Forecast Revisions</i>		
	Fed Funds	Eurodollars (1Y)	GDP
Monetary Path Shock	-0.03 [-0.11, 0.01]	0.39 [-0.02, 1.13]	-0.68 [-1.14, -0.45]
Information Shock	0.04 [-0.04, 0.15]	0.94 [0.67, 1.46]	0.64 [0.22, 1.25]
Monetary Target Shock	1.33 [1.01, 2.08]	0.32 [-0.06, 0.65]	-0.17 [-0.44, 0.03]

Panel B: Relative Variance in Regime 2,  $\Sigma_2 \Sigma_1^{-1}$

Monetary Path Shock	Information Shock	Monetary Target Shock
1.65	0.65	0.05
[0.79, 2.42]	[0.36, 1.57]	[0.02, 0.27]

This table shows the estimates of the system in equation (1.12), where the vector of observable variables has been expanded to include the HF surprise component of the Federal Funds rate announcement, as in equation (1.20). Otherwise, everything is identical to Table 1.3, whose note can be referenced for more detail.

Figure 1.14: Federal Funds Rate Surprise and the Target Shock



The top panel shows the 30-minute change in the price of the current-month Federal Funds rate future, scaled by the percentage of days left in the month (see, e.g., Kuttner (2001)), in basis points. The bottom panel shows the identified monetary policy target shock described in Section 1.5.1 standardized to have unit standard deviation over the plotted sample. The first two shaded areas present the regimes used for identification, described in Section 1.3.2. The third, gray shaded area, highlights observations that are dropped for most analyses, discussed in Section 1.2.4.

an effect on the surprise component of current-month interest rates.

The bottom row of panel A of Table 1.5 describes what I call a “monetary target shock.” This shock creates a positive comovement in short- and longer-term interest-rate forecast revisions, though the response is much stronger for short-term rates. GDP expectations are revised in the opposite direction, consistent with a theoretical monetary policy shock. Panel B shows that after the introduction of forward guidance, monetary target shocks essentially disappeared. With a few exceptions in 2007–2008, this remained true for the rest of the sample.

Figure 1.14 plots the Federal Funds rate surprise (top panel) and the estimated Federal Funds

target shock (bottom panel). The two series are fairly similar, consistent with the first column of panel A. This suggests that early estimates of monetary shocks were more closely aligned with exogenous monetary policy. The volatility of the two series also drops substantially in 2003, and never returns to its pre-2003 level. This disappearance causes a power problem: The estimated effects are far too noisy to draw any clear conclusions.


What drives this disappearance? Intuitively, once the Fed begins to communicate about future interest rates, it allows markets to forecast shorter-term interest rates better. Having received a signal about the Fed's time- $t$  interest-rate decision at an earlier date, there is little room for markets to be surprised by the time- $t$  interest-rate decision when it is announced. In Appendix B.2, I extend the illustrative model of Section 1.3.1 to a dynamic setting to show this point formally. Unlike the static illustrative model, in the dynamic model a perceived monetary policy target shock is the difference between the true shock revealed at time  $t$  and the pre-announced shock for time  $t$  made at meeting  $t - 1$ . The variance of target shocks is therefore unambiguously decreasing in the clarity of forward guidance. In the extreme case that the time- $t$  shock is revealed fully at time  $t - 1$ , there are no perceived monetary shocks.

Ramey (2016) suggests that monetary shocks are now rare because the Fed conducts policy more systematically and concludes that this is “bad news for econometric identification.” My findings corroborate the conclusion that (certain) monetary shocks are rare, but suggest an alternative mechanism and a different conclusion. The results in this section suggest that true monetary shocks—purely exogenous current-meeting interest rate surprises—*are* rare, as Ramey suggests. This conclusion may only be an artifact of the data that underlie my shocks: data on expectations.<sup>46</sup> With my data, these target-type monetary shocks largely disappear in response to a different mechanism: The Fed started announcing these shocks in advance in 2003, so their disappearance did not necessarily arise because the Fed has become more systematic. Few would argue, however, that over the last 50 years the Fed has become more systematic, so Ramey's mechanism is also likely

---

<sup>46</sup>As mentioned earlier in the paper, my shocks—alongside the shocks of any paper in the high-frequency literature—can only identify *perceived* shocks, since they are based on agents' perceptions (their expectations).

Figure 1.15: High-frequency Text-based Inflation Expectations



The left panel plots the average Blue Chip forecast revision of CPI inflation over the next three forecast horizons (analogous to  $\bar{Y}_{t|t+1} - \bar{Y}_{t|t+1}$  from equation (1.3)) against my HF measure of inflation forecast revisions,  $\hat{\pi}_t$ . The middle panel shows the relationship between my HF measure of inflation ( $\hat{\pi}_t$ ) and output forecast revisions ( $\hat{y}_t$ ). The final panel shows the relationship between  $\hat{\pi}_t$  and the information shock,  $\xi_t^I$ , identified in Section 1.3.4. Aside from Blue Chip forecast revisions, all other variables are standardized to have unit variance over the plotted sample.

at play.

On the other hand, my results suggest that monetary path shocks—commitments to deviate from the policy rule’s future prescriptions—are alive and well. These shocks and their effects only looked small because existing measures were swamped by information effects. Thus, rather than conclude that the Fed has put monetary econometricians out of business, I argue that finding new ways to measure the complexity of the Fed’s communications policies should keep them in business for the foreseeable future.

### 1.5.2 Supply vs. Demand Factors

The last question I ask in this paper is: What type of shocks does the public learn about when the Fed makes a policy announcement. This is an important question from a theoretic perspective: Jia (2020) shows that optimal communications policy depends on the nature of the underlying shock (in particular, cost-push (supply) vs. natural-rate (demand) shocks). To answer this question, I construct a measure of high-frequency text-based inflation expectations following the construction in Section 1.2.1. The only difference between the inflation and output index is in the

topic word list,  $\mathcal{Y}$ , which for inflation becomes

$$\mathcal{Y} = \{\text{inflat}, \text{price}, \text{oil\_prices}, \text{inflationari}, \text{deflat}\},$$

where I flip the sign of increasing and decreasing measures of `deflat`.<sup>47</sup> As with output, my measure of high-frequency inflation forecast revisions is the unpredictable component of post-meeting inflation directionality vis-à-vis pre-meeting directionality. I plot that index,  $\hat{\pi}_t$ , in the left panel of Figure 1.15, against the summary statistic of CPI forecast revisions from the Blue Chip survey, analogous to the measure constructed in equation (1.3). The correlation coefficient of the two series is 0.26, with a robust standard error of 0.08. In table 1.6, I show that this measure provides a high-frequency measure of inflation expectations that is not contaminated by risk premia, in contrast to what can be estimated using TIPS markets. There, I show regressions of the form

$$x_{t+1} = \alpha + \beta \hat{\pi}_t + \rho x_t + e_t \quad (1.21)$$

where  $x_t$  is either the (monthly) measure of 10-year inflation expectations or 10-year inflation risk premia estimated by Haubrich et al. (2012), and  $\pi_t$  is either the newspaper-based measure, or the one-day change in 10-year breakeven inflation. The only significant  $\beta$  in the table is on the relationship between (risk-neutral) inflation expectations and the newspaper-based measure, suggesting its usefulness in separately identifying changes in inflation expectations.

In Section 1.4.1 I showed that the identified information shock induced a positive correlation between real GDP and inflation and in expectations thereof. The right panel of Figure 1.15 shows that this relationship holds in high frequency. Specifically, the figure shows a scatter plot of the expansionary information shock (which, recall, increases output expectations) and  $\hat{\pi}_t$ . The two have a positive correlation of 0.13, with a robust standard error of 0.06.

The evidence thus leans against the notion that markets learn about supply-type factors from

---

<sup>47</sup>The bigram `interest_rates` was also included as similar to my “seed” words of `inflation` and `prices`, but I removed it.

Table 1.6: Comparison of Newspaper-Based Inflation Expectations with Breakeven Inflation

	Inflation	RP	Inflation	RP
Lagged LHS	0.963 (0.0180)	0.726 (0.0254)	0.971 (0.0180)	0.729 (0.0257)
Newspaper Measure	0.0162 (0.00686)	0.00322 (0.00311)		
TIPS Based			0.201 (0.242)	0.0420 (0.0841)
Observations	164	164	164	164
$R^2$	0.952	0.675	0.951	0.673

This table shows estimates of equation (1.21), described in the text. Heteroskedasticity-robust standard errors are shown in parentheses. The “Inflation” and “RP” columns show, respectively, represent the monthly measure of 10-year inflation expectations and 10-year inflation risk premia estimated by Haubrich et al. (2012) (these are measured at the beginning of each month,  $t$ ). The sample includes all regularly-scheduled FOMC meetings from May 1999 until October 2019.

policy announcements. That does not imply that markets learn *nothing*, only that information about supply factors is less prevalent. My results suggest, additionally, that it may be difficult to identify such a shock. The middle panel of Figure 1.15 shows that a fairly strong positive relationship exists between output and inflation forecast revisions ( $\hat{y}_t$  and  $\hat{\pi}_t$ ). The two series have a positive correlation of 0.27 (s.e. 0.10). That relationship is fairly consistent along my entire sample. Thus—at least given the heteroskedasticity-based approach to identification—there does not appear to be a natural regime that would help to separately identify a shock revealing supply-type factors from a shock revealing demand-type factors. When carrying out the estimation on a quad-variate system that includes high-frequency revisions in short-term rates, long-term rates, output, and inflation, the two non-monetary shocks have statistically equal variances across the two regimes. Identification by heteroskedasticity requires that the variance of different shocks differ across the two regimes, which means that the two information shocks are not identified. This is not surprising; nothing about the regime shift in 2003 suggests that markets could learn more about supply or demand shocks after 2003. That regime shift instead only allowed markets to learn relatively more about longer-term interest rates than about macroeconomic information in general.

Interestingly, the conclusion that markets learn little about supply-side factors from Fed an-

nouncements is consistent with the optimal policy prescriptions of Jia (2020). Jia highlights—in a much more realistic model than mine—the fact that optimal communications policy by the Fed reveals information about demand shocks, but obfuscates information about supply shocks (cost-push shocks, to be precise). Purposeful or not, my results suggest that the Fed’s communication policies are in line with the optimal communications policy.

## 1.6 Conclusion

Estimating the macroeconomic effects of monetary policy is notoriously difficult, because interest rates are so highly endogenous with respect to macroeconomic variables. Despite several approaches in the literature aimed at identifying exogenous changes in interest rates (i.e., monetary policy shocks) using high-frequency data, the identification of a series that has theoretically consistent effects, without relying on a particular set of controls, has proven elusive.

In this paper I provide an estimate of monetary shocks that is free of the “information effects” that have been posited to plague previous high-frequency estimates.<sup>48</sup> To separately identify two shocks—an information shock and a monetary shock—I study two data series: high-frequency interest-rate and GDP forecast revisions. Not able to find the latter variable in the literature, I constructed one using newspaper articles written about Fed policy meetings.

The notion that two series do not identify two shocks, without additional identifying assumptions, is an old one in economics. I present a model of expectation revisions around Fed announcements, and show what assumptions have been made (largely implicitly), by those who have previously sought to identify monetary shocks in high frequency. Like zero restrictions in structural VARs, these are generally not supported by macroeconomic theory. My assumptions, identification by heteroskedasticity, are a bit more involved, but much less restrictive and better suited to this context. The introduction of forward guidance by the Fed, a major regime change in Fed communications, drastically changed the shocks that markets could learn about from Fed policy announcements. Identification by heteroskedasticity provides a tool to turn this regime change into

---

<sup>48</sup>My results confirm that about 80% of a commonly used high frequency shock are made up of information effects.

structural identification.

In terms of empirical findings, my monetary shock has macroeconomic effects that are consistent with standard New Keynesian models. I contrast my estimates with other leading alternatives, and show that my results do not depend on having particular controls in my regressions. This lends credibility to my estimated shocks and speaks to their portability. Given the widespread use of monetary shocks in the empirical macroeconomics literature, a credible and portable series that can be carried forward (i.e., not limited by zero lower bound constraints) is greatly needed.

My results reveal additional information on the nature of monetary policy shocks. It appears that unanticipated shocks to the current-month Federal Funds rate have become a thing of the past, largely ending in 2003, with the introduction of forward guidance. Instead, as a result of the Fed's clearer signals, markets were able to more completely update their longer-term interest-rate expectations, leading to an increase in the size of "path" shocks. My model highlights how the introduction of forward guidance can lead to both smaller target shocks and larger path shocks.

Finally, I find that the effects of Fed information shocks are essentially equal in absolute value to the effects of monetary policy shocks. My results suggest that Fed announcements primarily reveal information about aggregate demand. This is in keeping with the prior observation (Jia, 2020) that optimal central bank information provision policy primarily reveals information about demand shocks. Further research will serve to clarify why the Fed appears to play such an important role in providing macroeconomic information, and how such communications can be further optimized.<sup>49</sup>

---

<sup>49</sup>The large macroeconomic effects that I estimate require reconciliation with empirical work that shows that households know very little about what the Fed does and says (Binder, 2017; Coibion et al., Forthcoming).

## Chapter 2: The Regressive Nature of the U.S. Tariff Code: Origins and Implications

### 2.1 Introduction

In trade models, tariffs are typically represented as a simple  $\tau$ —a wedge between the buyer’s and seller’s price of traded goods.<sup>1</sup> In reality, the U.S. tariff schedule is tremendously complex—today’s schedule comprises 4,394 pages of tariffs on 19,347 varieties of goods. Adding to this complexity, tariffs are not set uniformly across goods or industries, but are set at the *variety*-level. The monolithic nature of the tariff code nearly guarantees that individual tariff rates are not re-optimized each period, but are hysteretic: in many cases, tariff rates reflect tariff policies of bygone eras that were set with era-specific objectives in mind. Estimating the aggregate and distributional consequences of tariff policy requires taking this complexity into account. Documenting the existence and origins of tariff rates at the variety level is challenging, however, given the complexity of the tariff schedules and lack of digital records.

In this paper, we document a fact in today’s tariff schedules that is emblematic of the hysteretic nature of the U.S. tariff code: tariffs are regressive in nature, with tariffs systematically higher on low-value versions of goods relative to their high-value counterparts. For example, the tariff on a \$400 handbag made of reptile leather is 5.3 percent, while the tariff on an \$8 plastic-sided handbag is 16 percent. To this point, this pattern has only been noted anecdotally. Most recently, Furman et al. (2017) describe a number of consumption goods categories for which unit values (a

---

<sup>1</sup>This chapter is co-authored with Lydia Cox. We thank Pol Antràs, Elhanan Helpman, Réka Juhász, Marc Melitz, Emi Nakamura, Nathan Nunn, Ken Rogoff, and David Weinstein for helpful comments, as well as seminar participants in the Graduate Student Workshops in International Economics and Economic History at Harvard and the International Trade Colloquium at Columbia. Blaine Helleloid, Paulina Parsons, and Yi Yao provided excellent research assistance. We gratefully acknowledge financial support from the Harvard University Lab for Economic Applications and Policy (LEAP), the Program for Economic Research at Columbia, and the Alfred P. Sloan Foundation Pre-doctoral Fellowship in Behavioral Macroeconomics, awarded through the NBER.

proxy for prices) are negatively correlated with statutory tariff rates.<sup>2</sup> We show that this pattern is present across the entire tariff schedule, and is especially pervasive among consumer goods. Using newly-digitized tariff schedules going back to the 1930 Smoot-Hawley Act, we show that regressivity originated in the 1930s and 40s and has persisted over time. This pattern is one that arises from “within good” tariff variation—i.e., tariffs on lower and higher end varieties of goods are set differently. Our data allow us to show that, across the entire tariff schedule, most within good variation in tariffs originated in the few decades following Smoot-Hawley, and persisted to today. While regressivity may have been a cost of policy objectives of the 1930s and 40s, we show that, today, this pattern substantially alters the implied distributional gains from trade.

Regressive tariffs are present throughout the modern U.S. tariff code. We compare tariff rates among *varieties* of narrowly defined *goods*, and find that for about 60 percent of the goods with any tariff variation, lower cost varieties face a higher tariff than the higher cost varieties of the same good. We call the goods that fit this pattern *regressive goods*. For consumer goods, over 70 percent of goods with any tariff variation are regressive, and low cost varieties face a tariff rate that is, on average, over a percentage point higher than the rates on the high cost varieties. Among regressive goods, this difference in average tariff rates is over 5 percentage points. Compounding the consequences of this discrepancy is the fact that regressivity is not a new phenomenon. Over the last 30 years, relative tariff rates on low- versus high-value varieties have been remarkably persistent, suggesting that the divergence in rates between these two groups emerged long ago, and raising the question of whether relative rates today reflect current conditions or conditions of the past. Though the availability of and complexity of data make it difficult to perform systematic analysis of tariff rates over the near century between 1930 and today, we come at the question from many angles and find a similar story: as a whole, tariff rates have come down over time, but relative tariff rates among varieties of the same good are largely functions of trade negotiations of the past.

---

<sup>2</sup>Specifically, they note that among a subset of consumer goods “that have a calculated correlation between tariffs and prices, the tariffs scale decreases in prices for roughly half, and the list of categories with strong negative correlations is extremely wide-ranging.” Also see Gresser (2003).

Through detailed case studies of three goods—fishing reels, forks, and bicycles—we find that the divergence between tariff rates on low-value versus high-value varieties of many goods originated in the 1930s and 40s and has persisted over time. That divergence arose largely from two policy priorities: concessions to key trading partners and protection for domestic industries. In many cases, tariff rates were reduced as concessions to important trading partners, who tended to be advanced economies producing higher-end varieties of goods. Countries that were considered core participants of General Agreement on Tariffs and Trade (GATT) negotiations in the mid-twentieth century were 30 percent more likely to export high unit-value varieties. In other cases, tariff rates remained elevated on low-value varieties in order to protect a domestic industry from import competition, when that domestic industry specialized in production of the low-value variety. Our evidence suggests, however, that today these policy priorities are of much less relevance. On concessions, core GATT participants export much less of high-value products than they did in the first half of the twentieth century. On protections, the U.S. has largely shifted away from exporting low-value varieties. It is in this sense that we say that U.S. tariffs exhibit hysteresis: much of their variation today shows that consumers would save a little over \$4 billion on imported goods (via savings on tariff revenues) if the regressive pattern were eliminated, \$3 billion of which would come from a subset of consumer goods, totaling about 1.5% of their imports. These savings would likely not be shared equally across the income distribution. Using the sufficient statistics approach from Borusyak and Jaravel (2021) and a reduced form exercise, we show that regressivity amplifies the pro-poor effects of trade liberalization. In a baseline counterfactual tariff schedule in which we lower tariffs uniformly by 5 percentage points across all varieties, the average welfare gain for an individual earning \$20,000 per year is about two times higher than an individual earning \$100,000 per year. When we additionally eliminate the regressivity in tariff rates, the lower-income individual gains about three times as much. These relative gains are even more pronounced among consumer goods.

This paper is related to several strands of literature: first, the small literature on classification in trade. Related to our finding that tariff rates were often set in order to appease certain trading

partners, Grant (2021) argues that the entire system used to classify traded goods in the United States is determined endogenously, as the policymaker weighs the benefits of better policy targeting against the costs of more complex classification schemes. In earlier work, Gowa and Hicks (2018), Gulotty (2018), and Tavares (2006) find that tariff lines are often split or reclassified to accommodate political goals. Our work provides evidence of this type of endogeneity, and its economic consequences. We also put these previous findings in a dynamic setting: not only is classification endogenous, but it can be hysteretic as well. Endogenous classification decisions are not necessarily revisited and re-optimized in each period, but can persist even as the economic landscape shifts.

We also contribute to the literature that seeks to understand the political economy of the world trade negotiations that took place throughout the 20th century and their subsequent effects. Bagwell et al. (2020) analyze recently declassified tariff bargaining data from the Torquay Round of the GATT and document that negotiations were characterized by a lack of strategic behavior among participants and an important multilateral element to bargaining. Bown and Irwin (2015) study tariff levels in the late 1940s and early 1950s and find that tariffs fell by relatively more in the early rounds of the GATT for a core group of GATT participants—the United States, United Kingdom, Canada, and Australia—than they did for many other important countries (including other non-core GATT participants). Our work shows the persistent economic consequences of these negotiations.

We also contribute to a growing literature on the distributional effects of trade. Estimating the gains from international trade has been a central and long-standing question for policymakers and researchers. A more recent strand of the literature has acknowledged that the gains from trade may be distributed unevenly because individuals across the income distribution differ widely in the goods that they consume. Using data on aggregate expenditures and a nonhomothetic demand system, Fajgelbaum and Khandelwal (2016) find that trade favors individuals at the lower end of the income distribution, who tend to concentrate spending in more traded sectors. He (2018) also finds that trade has pro-poor effects and reduces real-wage inequality, and Hottman and Monarch (2020) find that, due to non-homotheticities, low-income consumers experienced more import price

inflation than high-income consumers between 1998 and 2014. On the other hand, Borusyak and Jaravel (2021) find that the purchasing-power gains from lower trade costs are distributionally neutral. While we do not present a fully specified structural model, as in the aforementioned studies, our reduced form estimates align with the studies finding that trade has pro-poor effects. We find that individuals with lower income tend to consume relatively more of low-value varieties. Taking this detailed regressivity into account has an effect over and above previous findings, shifting the distribution of gains from trade liberalization even more toward the left tail.

The rest of the paper proceeds as follows: In Section 2.2, we define regressivity and document that it exists throughout the modern-day tariff code. In Section 2.3, we perform detailed case studies on a few consumer goods in order to illustrate when and why regressivity emerged. In Section 2.4, we describe the newly digitized dataset on legislated tariff rates since 1930 that we rely on for our historical analysis. We use this data, in Section 2.5, in order to bolster the case studies with some more systematic evidence. In Section 2.6, we consider the welfare implications of regressivity, then we conclude in Section 2.7.

## **2.2 Documenting the Pattern**

In this section, we document the pattern that is this paper’s focus: tariffs are higher on low-value varieties of goods. We start by defining what a “good” is in the context of the Harmonized Tariff Schedule (HTS) in Section 2.2.1, and present summary statistics of those goods in Section 2.2.2. In Section 2.2.3, we define and present metrics by which to summarize the tariff differences on varieties of each good. In Section 2.2.4, we present alternative definitions and robustness checks. Our baseline analysis is performed on the 2017 U.S. tariff schedules; in Section 2.2.5, we ask whether the pattern has held over time (over the last 30 years) and in other countries (the EU). Our analysis reveals that the regressive pattern is not a new phenomenon, nor is it unique to the United States. These findings motivate the historical analysis presented in Sections 2.3 and 2.5.

### 2.2.1 Defining a “Good”

Documenting the regressive pattern in the tariff code requires a careful comparison of tariff rates and unit values among varieties of narrowly-defined goods. The World Customs Organization classifies every internationally-traded good with a six-digit Harmonized Commodity Description and Coding Systems (HS) numeric code. Countries are then permitted to provide more detailed classifications by adding digits to each good’s HS code. In the United States, these detailed classifications take the form of ten-digit Harmonized Tariff Schedule (HTS) codes, which are administered by the U.S. International Trade Commission and enforced by the U.S. Customs Bureau. Following convention, the term “HS code” is used herein to describe HTS or HS codes, unless the distinction is important. Eight-digit codes define a good’s “U.S. rate line,” and ten-digit codes describe its “non-legal statistical reporting category.”<sup>3</sup> Accompanying these numeric codes are text descriptions of each good. Because an important part of our analysis is ensuring that we are comparing the tariff rates on different varieties of the “same good,” we rely on the text descriptions to help define “goods” within the HTS. In robustness exercises, we also shows that our findings persist when using an approach to defining goods that does not rely on the text descriptions.

The best way to see how we use the text descriptions of each variety to define goods is by example. Table 2.1 presents all of the varieties (ten-digit products) within the four-digit HS code 8215. The format of the HTS is such that product descriptions get more and more detailed as additional digits are added to the HS code. Unfortunately, because of the vast coverage and complexity of the HTS, there is no simple rule that establishes an HS-digit level  $N$  such that all HS- $N$  codes define goods, and HS- $M$  codes (for  $M > N$ ) are varieties of that good. That fact is apparent here, with Spoons falling into two different HS-6 codes (8215.91 and 8215.99). Additionally, in some

---

<sup>3</sup>For the vast majority of products, tariffs are set at the eight-digit level. Currently, the exceptions are five eight-digit HTS codes corresponding to copper ores, lead ores, silver ores, other precious-metal ores, and “ash and residues containing mainly zinc.” See [https://www.usitc.gov/tariff\\_affairs/about\\_hts.htm](https://www.usitc.gov/tariff_affairs/about_hts.htm) for the naming of eight- and ten-digit codes, and [https://pubapps2.usitc.gov/tariff/readme\\_hts.jsp](https://pubapps2.usitc.gov/tariff/readme_hts.jsp), which notes that eight-digit goods with more than one tariff rate (i.e., eight-digit codes with different tariffs at the ten-digit level) have NA as the unit of quantity. Within eight-digit codes with NA units of quantity, only the five mentioned above have ten-digit tariff rates that are not uniform within the eight-digit code.

Table 2.1: 2019 HTS: All Varieties with HTS Code 8215

HTS Code	Description	Noun Grouping
8215	Spoons, forks, ladles, skimmers, cake-servers, fish-knives, butter-knives, sugar tongs and similar kitchen or tableware; and base metal parts thereof:	
8215.10.00.00	Sets of assorted articles containing at least one article plated with precious metal	sets
8215.20.00.00	Other sets of assorted articles	sets
	Other:	
8215.91	Plated with precious metal:	
8215.91.30.00	Forks	forks
8215.91.60.00	Spoons and ladles	spoons
8215.91.90.00	Other (including parts)	–
8215.99	Other:	
	Forks:	
	With stainless steel handles:	
	With handles containing nickel or containing over 10 percent by weight of manganese:	
8215.99.01.00	Valued under 25¢ each, not over 25.9 cm in overall length	forks forks
8215.99.05.00	Other	forks
	Other:	
8215.99.10.00	Valued under 25¢; each	forks
8215.99.15.00	Other	forks
8215.99.20.00	With rubber or plastic handles	forks
	Other:	
8215.99.22.00	Without their handles	forks
	Other:	
8215.99.24.00	Table forks (including table serving forks) and barbecue forks with wooden handles	table forks
8215.99.26.00	Other	table forks
	Spoons and ladles:	
	With stainless steel handles:	
8215.99.30.00	Spoons valued under 25¢each	spoons
8215.99.35.00	Other	spoons
8215.99.40	With base metal (except stainless steel) or nonmetal handles	–
8215.99.40.30	Tablespoons and table ladles	tablespoons
8215.99.40.60	Other	tablespoons
8215.99.45.00	Other	–
8215.99.50.00	Other (including parts)	–

NOTE. This table presents an example of the layout of the HTS in 2019, for the section of the HTS that describes tariffs on goods with HS-4 code 8215. The indentations replicate what is in the HTS, and the column “Noun Grouping” shows the relevant noun (or pair of nouns) that define the several goods within the HS-4 8215—see the text for more details.

cases, varieties of two different goods might be classified under the same HS-*N* code. For example, 8215.99 contains forks and spoons, with forks running from 8215.99.01 through 8215.99.26, and spoons starting thereafter with 8215.99.30.

Ameliorating these complications is the layout of the HTS, which uses different levels of indentation to highlight relevant splits of HS codes. Returning to the example in HS 8215.99, the fork/spoon demarcation becomes clear once we leverage these indentations: all the forks lie under the “Forks” heading (or have descriptions beginning with “forks”), and spoons lie under the “Spoons and ladles” heading (or have descriptions beginning with “Spoons”). Each additional level of indentation introduces additional descriptive text for the products subsumed by it (that is, by varieties that are more-indented). When this descriptive text or product description begins with one or two *nouns*, and no nouns are used to start the description of any more-indented varieties, then that the noun becomes the definition of a good for that product and all products subsumed by that indentation.<sup>4</sup> We find the part of speech of each word using the Penn Treebank project, described in Marcus et al. (1994), which, while not particularly new, is still the natural language processing industry standard.<sup>5</sup> To minimize the risk that different nouns could be used describe different goods (for example, “fish” could be used to describe the animals or “fish hooks”), only noun-groupings within HS-4 groupings are considered “goods.”<sup>6</sup> So, in the example of table 2.1, the goods introduced are 8215-sets, 8215-spoons, 8215-table forks, 8215-tablespoons, and 8215-forks.<sup>7</sup>

---

<sup>4</sup>We also consider adjective-noun pairs, which (anecdotally) does not split goods into varieties.

<sup>5</sup>This tagger considers not only each word to be classified, but also its surrounding words. So, for example, in “I went fishing,” fishing is a verb, while in “fishing reels,” fishing is a noun.

<sup>6</sup>The level of detail and margin of delineation for goods described in the HTS reflect trade patterns. So, for example, while meat of horses, asses, mules, or hinnies are all classified by the four-digit HS-code 0205, with no further delineation, the meat of lamb is broken down into ten-digit codes depending on the body part for bone-in meat, and eight-digit codes for boneless meat. This, and countless other examples, support the claim that defining a good based on its HTS code alone leads to definitions of goods that are inconsistent with what a human reader would define. However, at a high-enough level of aggregation the HS codes are useful for separating types of goods.

<sup>7</sup>From a text-processing perspective, this technique produced far more reasonable results than other more-sophisticated natural language processing tools. Perhaps the easiest way to see why an untrained algorithm would have difficulty in defining a good is to consider again the forks example. Suppose that the four varieties were silver forks, silver spoons, steel forks, and steel spoons. Without any additional information, algorithms that measured the similarity of each string (like cosine similarities under Latent Semantic Analysis, or Levenshtein distances) have no

Table 2.2: Good-Level Summary Statistics

	All	Consumer	Intermediate	Capital	Other	Low-Tech. Consumer
Number of Goods	2,987	956	1,649	387	46	474
Number of Varieties	12,487	4,649	6,428	1,222	188	3,030
Imports (\$, billions)	1,334	339	543	306	147	209
Import-weighted Tariff	0.019	0.051	0.010	0.004	0.009	0.076
Fraction w/ tariff variation (by count)	0.644	0.689	0.660	0.457	0.391	0.827
Fraction w/ tariff variation (by imports)	0.549	0.620	0.571	0.286	0.847	0.699
Fraction with $\leq 2$ varieties	0.523	0.456	0.560	0.584	0.652	0.340
$\leq 4$ varieties	0.779	0.722	0.800	0.860	0.826	0.584
$\leq 10$ varieties	0.934	0.903	0.945	0.982	0.935	0.831

NOTE. This table presents summary statistics for different slices of the 2017 U.S. tariff schedules. See the text for details on the sample selection.

### 2.2.2 Coverage of the Tariff Schedule

Our main analysis focuses on the subset of total imports into the United States that can be reasonably classified into “goods.” As such, we drop a few classes of imported varieties. First, we drop two chapters of the schedules (98 and 99) that reflect special and temporary tariffs. Second, we drop goods with no listed unit of quantity, since our analysis requires us to define the unit value of imported varieties. Third, we drop varieties that were not classified into a “good” using the algorithm described above. The vast majority of these cases arose from varieties that were described as “other” in the tariff schedules, as in the example of Table 2.1. Finally, we only analyze goods with at least one variety, since our focus is on within-good tariff variation. Overall, our data cover \$1.3 billion of \$2.3 billion in total U.S. imports of goods in 2017. The largest losses in coverage of imports come from the second (no listed unit of quantity) and final (only one variety) steps, which represent about \$650 billion of imports.

Table 2.2 presents summary statistics of the goods and varieties covered by our analysis. Since our objective is to demonstrate the pervasiveness of the regressive pattern across the tariff schedule, we present statistics for the full dataset, as well as for a few key subsets. The first three subsets—Consumer, Intermediate, and Capital goods—are categorized using the UN’s Broad Economic Category (BEC) classification of 6-digit HS codes.<sup>8</sup> The final column represents a category of

---

way to know that spoon/fork is the right delineation, rather than steel/silver. The algorithm defined above leverages the fact that human beings created the layout of the HTS.

<sup>8</sup>The BEC categorizes transportable goods according to their end use. The classification is more de-

goods that we call low-tech consumer goods or LTCGs. This category consists of all (6-digit HS) varieties that are classified as consumption varieties under the BEC, but exclude food, electronics, and vehicles.<sup>9</sup> What remains are the types of household products that consumers purchase at a regular frequency, akin to non-durable consumption.

There are a few important takeaways from these summary statistics that illuminate the heterogeneity that exists in tariff rates across the tariff code. First, though tariff rates are often thought of or modeled as being uniform across good or industries, the fifth row of Table 2.2 shows that there is a substantial amount of variation in tariff rates at very disaggregated levels. Whether by count or by value, the majority of goods in each category have within-good variation in tariff rates. This is especially true for LTCGs, in which 83 percent of LTCG goods have within-good tariff variation. This may be in part because LTCGs also have a higher average number of varieties-per-good than the other categories, suggesting that these goods have received more attention from policymakers (in the framework of e.g. Grant (2021)).<sup>10</sup> Lastly, consistent with the conventional wisdom that the United States has relatively low tariffs on average, the average (import-weighted) tariff rate across all goods is only 1.9 percent. Aggregation, however, masks substantial heterogeneity in tariff rates in different categories. Consumer goods and LTCGs have much higher average tariffs of 5.1 and 7.6 percent, respectively. In fact, LTCGs made up only about 16 percent of imports in 2017, but accounted for 63 percent of total import duties.

### 2.2.3 Regressivity Across the Tariff Code

To evaluate the pro/regressivity of tariff rates, we start by sorting (HS10) varieties into within-good quantiles based on their unit-values. Within each good, we classify varieties in the top quan-


---

tailed than the three categories we present here, which are based on the categorization of BECs into three System of National Accounts categories, plus a residual. The details of both categorizations can be found in United Nations (2002). The crosswalk from BEC codes to 2017 HS codes is available from <https://unstats.un.org/unsd/trade/classifications/correspondence-tables.asp>.

<sup>9</sup>Specifically, we exclude varieties from sections I–IV of the HS (food), section XVI (machinery), and section XVII (transportation equipment).

<sup>10</sup>The consumer good with the most varieties is cheese, with 59 varieties, whose low-value variety include Velveeta, and whose high-value variety consists of Edam and Gouda cheeses.

Figure 2.1: Heterogeneity and Overall Presence of the Pattern



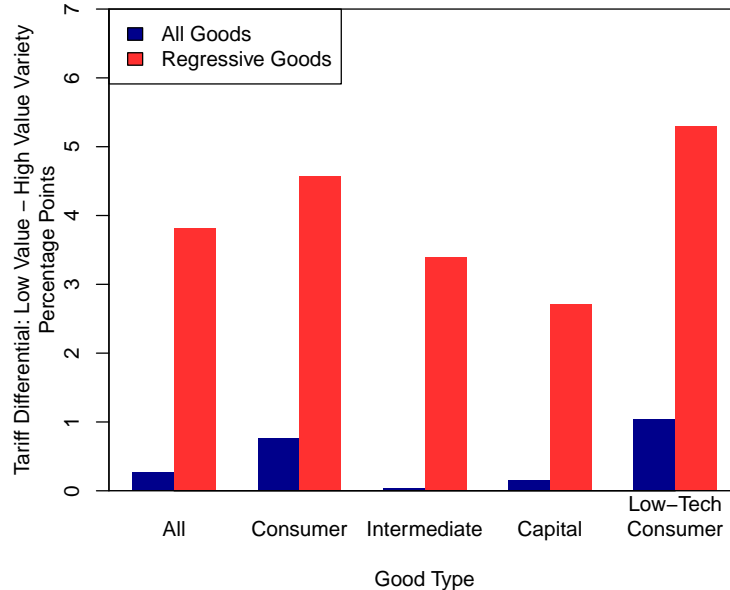
NOTE. This figure reports the share of regressive goods within each product category listed on the x-axis. Regressive goods are goods for which the average MFN ad valorem tariff on low-value varieties is greater than the corresponding average tariff on high-value varieties. 95% confidence intervals are shown with whiskers, computed using heteroskedasticity-robust standard errors. See Section 2.2.2 for a discussion of the categories and sample.

tile as “high-value” varieties, and those in the bottom quantile as “low-value” varieties. Our baseline approach splits goods at the median, though we show below that we obtain similar results when we split goods into quartiles.

We find that regressivity is pervasive across the tariff schedule. Figure 2.1 shows the share of goods for which the average tariff on low-value varieties is higher than the average tariff on high-value varieties. We compute this statistic for different industry groupings using ad-valorem equivalent most-favored-nation (MFN) rates.<sup>11</sup> The dark blue bar on the right shows this breakdown for all goods. In the aggregate, 60 percent of the goods with any tariff variation exhibit the regressive pattern. This share varies widely across industries (the gray bars) and industry groups (the light blue bars, which depict BEC-classified categories). Intermediate goods industries like

<sup>11</sup>Data on legislated tariff rates (i.e., ad-valorem rates) come from the U.S. International Trade Commission’s website and Feenstra et al. (2002). Data on trade flows and effective tariffs come from Schott (2008). The ad valorem equivalent rate is the sum of the ad valorem rate with the quotient of the specific rate and unit value.

Figure 2.2: Tariff Differentials between Low- and High-Value Goods



NOTE. The graph shows the average tariff differential between low- and high-value varieties among each category on the x-axis. The differential is constructed, for each good, as the difference between the average MFN ad valorem tariff on low-value varieties and the corresponding tariff on high-value varieties.

metals, plastic, and rubber are at the low end of regressivity, with fewer than half of goods in those industries exhibiting the pattern. Among consumer goods and LTCGs, on the other hand, almost three quarters of the goods that have any tariff variation are regressive goods.

Not only are tariffs on low-value goods higher than tariffs on high-value goods, but in many cases, the discrepancy in tariff rates within goods is substantial. Figure 2.2 shows the average tariff differential between low (below median) and high-value (above median) varieties in each category of goods. The dark blue bars show the differential for all goods and the red bars restrict the sample to only the “regressive” goods, or the goods that fit the regressive pattern. Averaging across all goods, tariff rates on low-value varieties are around 0.4 percentage points higher than tariff rates on high-value varieties. Among regressive goods, the differential is around 4 percentage points. Differentials are magnified in the LTCG category. Across all LTCGs, tariff rates on low-value varieties are 1.24 percentage points higher than the rates on their high-value counterparts, and among regressive LTCGs, the differential is 5.3 percentage points.

The fact that regressivity is so commonplace across the tariff code is a hint that the pattern is

economically significant, not just a harmless statistical quirk. The statistics presented above are based on unweighted averages of varieties and goods, which we focus on as our baseline due to the potential endogeneity of imports to tariff rates. In Appendix K, we show that the regressive pattern holds when we use import-weighted averages as well. More discussion of the economic significance of regressivity can be found in Section 2.6.

#### 2.2.4 Robustness: Alternative Definitions and Global Presence

There are other justifiable ways to define and calculate the statistics presented above. In Table 2.3, we show that alternative statistics and subsamples tell the same story. In the table, the baseline analysis of Figures 2.1 and 2.2 are shown in tabular form in the first, fifth, and sixth rows, respectively.

The shares of regressive goods, shown in the second through fourth rows, are quite stable across various alternative definition of a “regressive good.” The second row shows our estimates when we split goods into quartiles by unit value, rather than in half. Put differently, a good here is “regressive” if the average tariff on the bottom quartile goods is higher than the average tariff on the top quartile. In the third row, we alter the definition of a “good” from the text-based measure in our baseline, to a five-digit HS code/unit of quantity pair. The results are consistent with the baseline, showing robustness to the description-based definition of goods above. To address the fact that specific tariffs (a fixed duty per unit) can mechanically make a good look regressive, the fourth row uses only the ad valorem component of legislated tariffs.

Tariff differentials—shown in the fifth through ninth rows—are also quite stable across definitions. The fifth and sixth rows reproduce, respectively, the blue and red bars in Figure 2.2. While the use of only the ad-valorem component of rates reduces the tariff differential, shown in the seventh and eighth rows, in general, this is less pronounced for LTCGs, and does not meaningfully impact the fraction of regressive goods. In the ninth row, we redefine both goods, and the tariff differential, using effective tariffs (tariff revenue divided by imports) rather than statutory rates. This ensures that the rates account for 2017 sourcing patterns. The tariff differentials are essentially

unchanged.

In the bottom two rows, we show an alternative metric for defining regressivity: the average correlation of unit-values and tariff rates within a good. The same general takeaway holds here: tariffs tend to be lower on varieties with higher unit values. Among regressive goods, this correlation is quite pronounced, at -0.72 for all goods. Taken together, this evidence suggests that the pattern of regressivity is robust to our choice of data and definitions.

In appendix L, we show that the regressivity also pervades the tariff schedules of the European Union. In light of the historical analysis below, this will not be surprising. There, we highlight that much of this regressivity emerged as a result of early GATT negotiations, in which EU members (in 2017, i.e., including the U.K.) were the primary contributors. The presence of regressivity of a bloc of countries that (including the U.S.) make up about 30% of global imports, also has implications for the *producers* of low- vs. high-value varieties of goods. To the extent that low-value varieties are produced by lower-income countries (consistent with our historical analysis), the regressive pattern can have not only domestic, but also global distributional implications.

### 2.2.5 Regressivity is Not a New Phenomenon


The statistics presented thus far are representative of modern times, based off of the 2017 U.S. Tariff Schedule. Figure 2.3 shows, however, that regressivity is not a new pattern in the data. Figure 2.3a shows that not only have average tariff rates remained relatively flat between 1989 and the present, but the gap between rates on low- and high-value varieties has also been unwavering. This is even more pronounced for LTCGs, shown in Figure 2.3b. This recent stability raises the question of *when* this regressive pattern emerged, and, given the changes in the economic landscape and global sourcing patterns over the last 30 plus years, are these relative tariff rates still optimal, or are they remnants of a past era? In the next several sections of this paper, we seek to answer these questions.

Table 2.3: Establishing the Pattern: Robustness

	All	Low-Tech Consumer	Capital	Consumption	Intermediate
<b>Share Regressive</b>					
Average	0.60	0.71	0.67	0.71	0.47
(s.e.)	(0.01)	(0.01)	(0.03)	(0.01)	(0.01)
<b>Quartile Split</b>					
Average	0.60	0.71	0.66	0.71	0.47
(s.e.)	(0.01)	(0.01)	(0.03)	(0.01)	(0.01)
<b>HS5 Goods</b>					
Average	0.6	0.80	0.54	0.75	0.47
(s.e.)	(0.0)	(0.01)	(0.02)	(0.01)	(0.01)
<b>Ad-Valorem Only</b>					
Average	0.56	0.70	0.65	0.66	0.44
(s.e.)	(0.01)	(0.01)	(0.03)	(0.01)	(0.01)
<b>Tariff Differential</b>					
Average	0.41	1.24	0.30	1.07	0.03
(s.e.)	(0.11)	(0.28)	(0.15)	(0.26)	(0.13)
<b>Only Regressive</b>					
Average	3.99	5.30	2.7	5.00	3.40
(s.e.)	(0.20)	(0.36)	(0.4)	(0.35)	(0.25)
<b>Quartile Split</b>					
Average	0.54	1.45	0.28	1.49	0.00
(s.e.)	(0.13)	(0.33)	(0.15)	(0.30)	(0.12)
<b>Ad-Valorem Only</b>					
Average	0.16	1.11	0.24	0.53	-0.08
(s.e.)	(0.10)	(0.28)	(0.15)	(0.22)	(0.12)
<b>Effective Rates</b>					
Average	0.34	1.29	0.31	0.94	0.00
(s.e.)	(0.09)	(0.26)	(0.15)	(0.18)	(0.11)
<b>Unit-Value Tariff Correlation</b>					
Average	-0.09	-0.16	-0.09	-0.19	-0.01
(s.e.)	(0.02)	(0.04)	(0.09)	(0.03)	(0.03)
<b>Only Regressive</b>					
Average	-0.72	-0.56	-0.70	-0.66	-0.76
(s.e.)	(0.01)	(0.03)	(0.06)	(0.02)	(0.02)

NOTE. The table reports averages over different categories of products (across columns) of the different statistics and subsamples described in the text (rows). Heteroskedasticity-robust standard errors are reported in parentheses.

Figure 2.3: Legislated Tariffs, 1989-2017



NOTE. This figure shows the (simple) average legislated tariff rate for all varieties above and below the median unit value of a good. The left panel shows this calculation for all goods/varieties in the tariff schedule, while the right panel restricts the sample to LTCGs.

### 2.3 Origins of Regressivity: A Case Study Approach

The analysis presented in Section 2.2.5 suggested that the modern-day regressivity in the tariff schedules had its origins before the availability of digitized legislated tariff schedules. In order to understand when and why the pattern originated, we performed detailed case studies on three goods that exhibit regressivity today and are among a larger set of regressive consumer goods that we study more-systematically in Section 2.5: fishing reels, forks, and bicycles. Our detailed historical investigation, described in sections 2.3.1 and 2.3.2, suggests that the divergence in tariff rates between low- and high-value good dates back to the mid-1900s during the early rounds of the General Agreement on Tariffs and Trade (GATT). While each variety of good has its own history of tariff rate changes, the policy priorities driving the emergence of the regressive pattern tend to fall into one of two categories: concessions to important trading partners (lowering relative rates on high-value goods) or protection for domestic industries (raising relative rates on low-value goods). We discuss these in turn.

### 2.3.1 Concessions and the Case of Fishing Reels

First, trade negotiations and subsequent tariff concessions in the decades following the Smoot-Hawley tariff tended to be made with advanced economies that produced higher end varieties. It is widely accepted that, especially in its early years, the GATT was somewhat of a “rich man’s club.” Developed countries were the dominant players in negotiations, while developing countries were largely on the sidelines. Negotiations occurred according to a “principal supplier rule,” which dictated that agreements about tariff rates on a certain variety occur only with the principal supplier of that variety. Moreover, there was a sense of reciprocity in these negotiations—informal in some ways, but mandated by law in others<sup>12</sup>—such that countries would both ask for and grant concessions to trading partners. As a result, most of the “action” that occurred during the GATT rounds involved tariff reductions on goods supplied by developed countries.<sup>13</sup> Most developing countries were neither principal suppliers nor major importing markets, and so little was asked of them in terms of their own trade liberalization and the interests of their export markets were ill-represented (Bown, 2009). As an illustrative example, consider the following excerpt from the U.S. negotiating party at the Kennedy GATT round, which illustrates that even when developing countries were present at the negotiating table, it may have been difficult for their interests to be acknowledged due to differences in tariff nomenclature:

Due to the lack of precision in descriptions of the products submitted, substantial difficulties have been experienced in identifying in the United States tariff nomenclature the corresponding articles which were actually intended to be designated by the less-developed countries (United States Tariff Commission, 1965)

The well-known result of these negotiation dynamics is that while tariffs were lowered substan-

---

<sup>12</sup>According to Bown (2009), there is no article of the GATT 1947 that formally identifies reciprocity as a founding principle, however the articles that govern how countries are to renegotiate concessions (specifically, Articles XXVIII and XIX) if one party seeks to amend a bargain does contain language about reciprocity, suggesting that it played a role in how initial negotiations were carried out.

<sup>13</sup>And even more specifically, according to Bown and Irwin (2015), among a core group of GATT participants including the United States, United Kingdom, Canada, and Australia.

tially on a vast array of goods, trade barriers remained high in a number of markets that were of export interest to developing countries, like agriculture and clothing.

Somewhat more surprising, however, is that that this pattern appears to hold at the good-level as well. Anecdotal evidence suggests that the United States made tariff concessions primarily on *varieties* that were of interest to negotiating partners, distinguishing among varieties of the same good. According to a report of the Industry Sector Advisory Committee (ISAC), for example, about the Uruguay Round of tariff negotiations:

Since China was not a party to these negotiations, the inclination of U.S. negotiators was to avoid making tariff cuts on products if China would be the main beneficiary to prevent the free rider problem.<sup>14</sup>

In other words, at the variety level, efforts to lower trade barriers were largely determined by bilateral negotiations with principal suppliers, leading to within-good variation on tariff rates that favored negotiating partners.

The history of tariff reductions on fishing reels is a prime example of a case where tariff rates were lowered on expensive varieties to appease important trading partners. As shown in Figure 2.4, the tariff rate on fishing reels valued over \$2.70 and the tariff rate on fishing reels valued over \$8.45 was the same until around 1940—until then, these two varieties of fishing reels were not differentiated in the tariff schedule. The distinction between the high- and low-value fishing reels originated in the Anglo-American Trade Agreement of 1938. At that time, U.S. manufacturers dominated the domestic and foreign market for fishing reels—imports were less than one half of one percent of U.S. production.<sup>15</sup> Of the small number of imports that did come in, high-quality reels came from the United Kingdom and medium- to low-grade reels came from Japan and Germany. By creating a distinction in the tariff rates between high-quality reels (those valued above \$8.45) and low-quality reels (those valued below \$2.70)<sup>16</sup>, the United States was able to


---

<sup>14</sup>See United States International Trade Commission (1994).

<sup>15</sup>See United States Tariff Commission (1938) for a detailed analysis of this agreement.

<sup>16</sup>A third group was also created for reels in between \$2.70 and \$8.45.

Figure 2.4: History of Tariff Rates on Fishing Reels



NOTE. This figure shows the most-favored-nation tariff rate on fishing reels. Rates are from historical tariff schedules, written in the documents *U.S. Import Duties* between 1946 and 1963 (the first edition of which contains all rates since the 1930 Smoot-Hawley Act), *Tariff Schedules of the United States Annotated* between 1963 and 1989, and the HTS thereafter. Rates between 1913 and 1930 were extracted from Brossard (1930). The small numbers show instances in which the tariff-determining unit-values changed.

appease the United Kingdom in negotiations—lowering import barriers for the type of fishing reel that the United Kingdom produced—without creating a threat to domestic producers.<sup>17</sup> That is stated in no uncertain terms in the U.S. Tariff Commission’s 1939 Annual Report:

It is noted in the report that a large number of the duty reductions relate to articles, imports of which consist of grades or qualities either not produced in the United States or produced only in small amounts. Furthermore, it is pointed out that many new import classifications are established by the agreement either for the purpose of confining concessions to products supplied principally by the United Kingdom, or in order to limit duty reductions to the less competitive portions of the old classifications. United States Tariff Commission (1939)

<sup>17</sup>While the Anglo-American Trade Agreement of 1938 was a bilateral negotiation between the United States and the United Kingdom, the negotiated rates were actually applied as most-favored nation rates for both countries. For more on this treaty, and Cordell Hull’s strong roll in pushing through this agreement as a harbinger of world peace, see Schatz (1970).

Table 2.4: GATT Participant and Concentration of High-Value Exports (1974)

	(1)	(2)
	Above-Median Unit Value Export Share	Above-Median Unit Value Export Share
GATT Participant	0.230*** (0.052)	0.271** (0.087)
Observations	155	155

Standard errors in parentheses

Sample year: 1974.

Column (2) weighted by country's total imports.

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

NOTE. The core “GATT Participant” countries, described in the text, are Australia, Canada, the United Kingdom, the EEC (France, Germany, Italy, Denmark, Ireland), Benelux (Belgium, Luxembourg, and the Netherlands), and Japan. Heteroskedasticity-robust standard errors of the differences in averages are reported in parentheses.


Further concessions were made to the United Kingdom on expensive fishing reels during the GATT negotiations in Geneva in 1947, and though there were subsequent rate reductions for both varieties of reels during later rounds of GATT negotiations, the rates on the high- and low-value varieties never again equalized.

Evidence from trade flows is supportive of the notion that GATT participants were more likely to export high-unit value varieties across the board. To that end, we use the global trade flow data of Feenstra (1996) in order to compute, for each country, the share of exported varieties that are high-value varieties in 1974.<sup>18</sup> We then compute the average value of this export share in all GATT countries, and in all non-GATT countries. Table 2.4 shows the difference, demonstrating that the share of exports of all goods of above-median unit value varieties for core GATT participants was about 25 percent higher in 1974 relative to non-core GATT participants.<sup>19</sup> In other words, the primary participants in GATT negotiations were more likely to be exporting, and therefore negotiating on behalf of high-value varieties of goods.

<sup>18</sup>Our sample uses trade flows in 1974 because this is the earliest year for which we have bilateral trade data for all countries. We define a “good” here according to 5-digit TSUS codes. As was described in the context of table 2.3, that convention yields similar views on the presence of regressivity as our text-based approach.

<sup>19</sup>The list of “core” participants (Bown and Irwin, 2015) is Australia, Canada, United Kingdom, the EEC, and Japan.

Figure 2.5: History of Tariff Rates on Stainless Steel Forks



NOTE. This figure shows the most-favored-nation tariff rate on table forks with metal-plated handles. Rates are from historical tariff schedules, written in the documents *U.S. Import Duties* between 1946 and 1963 (the first edition of which contains all rates since the 1930 Smoot-Hawley Act), *Tariff Schedules of the United States Annotated* between 1963 and 1989, and the HTS thereafter.

### 2.3.2 Protection and the Cases of Forks and Bicycles

The other apparent driver of the divergence in tariff rates between high- and low-value goods in the mid-twentieth century is an age-old story: domestic industries lobbied for protection from cheap foreign imports. Article XIX of the GATT permitted contracting parties to escape GATT obligations and raise trade barriers to safeguard domestic producers that could prove serious injury would be caused by an increase of imports. In each round of GATT negotiations, a number of tariff lines were excluded from concessions for economic reasons, and in cases where the varieties in question were the low-value ones, regressivity was born: tariff rates on those varieties remained high, while negotiations lowered rates on the more expensive counterparts.

This protectionist motive played a prominent role in the case of forks—today forks of stainless steel valued under 25 cents have a 15.8 percent tariff, while those plated in precious metal enter freely. Shown in Figure 2.5, the divergence in these rates came about in the late 1950s. At that

time, there were 21 manufacturers of stainless steel table flatware in the United States. Most of these domestic firms were producing flatware of the low-value variety—valued between 8 and 25 cents per piece. At the same time, there was a sharp increase in imports of stainless steel flatware, and most of these imports were low-quality versions (valued around 14 cents per piece) coming from Japan. In search of protection from these cheap foreign imports, the domestic stainless steel flatware industry petitioned for, and received, protection under an escape clause tariff rate quota (TRQ) starting in 1959 for flatware valued under 25 cents per piece.<sup>20</sup> Since domestic production of precious metal-plated flatware was minor, the more expensive variety was not included in the escape clause protection<sup>21</sup>. As a result, during the Kennedy Round of GATT negotiations, the tariff rate on cheap stainless steel forks was excluded from concessions, remaining just under 20 percent, while the rate on silver-plated forks was lowered. Further rate reductions for silver-plated forks occurred in subsequent GATT rounds, but the rate on cheap stainless steel forks is roughly the same today as it was 60 years ago.

There are other cases, like the fork industry, where domestic industries specialized in production of the low-value variety of a good, leading to the divergence between tariff rates on high- versus low-value varieties in the face of import competition. The bicycle industry, for example, also petitioned for escape clause protection following concessions that were granted to the United Kingdom in 1939 and later multilateral concessions agreed to in the 1947 Geneva round of the GATT. In 1955, tariff rates on the cheapest three of four varieties of bicycles were raised from 15 to 22.5 percent. The tariff rate on the fourth and most expensive variety—lightweight bicycles (weighing less than 16.3 kilograms) was increased from 7.5 percent to only 11.25 percent<sup>22</sup>. The lower rate for lightweight bicycles was justified by the fact that virtually all lightweight bicycles were imported and did not directly compete with the most popular domestic model—the balloon

---

<sup>20</sup>Under the TRQ arrangement, imports of SSTF valued at less than 25 cents each were charged the concessionary tariff rate on the first 5.75 million dozen pieces imported. Above 5.75 million dozen pieces, imports were levied at a higher rate.

<sup>21</sup>USITC (1982)

<sup>22</sup>According to USITC (1982), the U.S. Tariff Commission actually recommended a rate of 22.5 percent for lightweight bicycles as well, but President Eisenhower halved this recommendation.

tire bicycle. Today the differentiation in rates remains. The two lightweight varieties of bicycles which have average unit values of \$200 to over \$1000 have a tariff rate of 5.5 percent or lower, while the rate on all other bicycle varieties that cost below \$150 on average is 11 percent.

### 2.3.3 Summary: Regressivity Came from Bygone Policy Priorities

If the two apparent drivers of tariff regressivity outlined in this section are indeed the causal forces, it is notable that the current pattern of tariff rates originated in a vastly different economic landscape than the one that exists today. U.S. trading partners in the mid-1900s and the types of goods that were imported from those partners have shifted dramatically over the last half-century, as has the composition of domestic industries. While concessions on high-quality fishing reels, for example, were made in the interest of the United Kingdom and world peace, in 2018 the U.K. accounted for less than 0.1 percent of U.S. imports of fishing reels valued over \$8.45. For forks, despite continued TRQ protection and an increase in domestic demand for stainless steel table flatware, the domestic stainless steel flatware industry has all but disappeared. A few domestic producers survived by diversifying production into higher-end flatware and other cookware, but today only one domestic manufacturer of stainless steel flatware remains: Sherrill Manufacturing (Liberty Tabletop). The retail price of a single stainless steel fork “Made in the USA” is \$3.49.<sup>23</sup> There are no domestic producers left being protected by the tariff on forks costing less than 25 cents per piece, yet the rate on those cheaper forks remains where it was in the early 1960s when the domestic industry was strong. In fact, tariff rates on many manufacturing industries were kept high during GATT negotiations, not to save domestic industries, but to cushion the blow of industries that were clearly already in decline. This was noted explicitly in the justification of the economic exception to tariff concessions given to the *Headwear of Fur Felt* industry during the Kennedy Round of the GATT:

Production and employment have declined persistently over the past decade in this in-

---

<sup>23</sup>See <https://www.nytimes.com/2015/10/22/business/smallbusiness/setting-a-made-in-the-usa-table-yes-do-stick-a-fork-in-it.html> for more on the state of the domestic flatware industry.

dustry. Imports now account for half of the domestic market. The few remaining firms and employees should be given additional time to complete the process of adjustment and diversification now underway. (United States Tariff Commission, 1965)

Though for most goods tariff rates have declined since the early GATT rounds, in many cases the divergence between rates on low- and high-value varieties has never been corrected. Rates on many low-value products remain high, even if those domestic industries no longer exist. In part, this is no surprise, as trade negotiations have become much more sophisticated over the years. Starting in the Kennedy Round of the GATT and continuing through the present day, trade agreement negotiations have been about more than just tariff rates, focusing on other important issues like provisions for intellectual property, environmental concerns, and labor protection. Moreover, given that firms, and not consumers, are the primary lobbying force for tariff rates, if there are no longer domestic firms in many of the industries in question in this analysis, it is not hard to believe that tariff rates on these products have been somewhat forgotten.

The next two sections are devoted to showing that the conclusions drawn from the case studies are indeed relevant when expanding the historical analysis in a more-systematic way. To that end, we bring in newly-digitized data on legislated tariffs since the 1930 Smoot-Hawley act and perform several cuts of that data that support the claims from this section.

## 2.4 New Data on Legislated Tariff Rates

We investigate the origins of the regressive pattern documented in Section 2.2 using newly-digitized data on legislated tariff rates going back to the Smoot-Hawley Act. Our data on legislated tariffs come primarily from the tariff schedules published by the U.S. International Trade Commission (formerly the U.S. Tariff Commission, TC).<sup>24</sup> We describe the documents in detail in

---

<sup>24</sup>Aside from the 1930 document, these documents contain only legislated tariff rates and lack trade-flow data. Until 1946 the Commerce Department (or preceding agencies) published the *Foreign Commerce and Navigation of the United States*, where we get data on rates in 1930. These documents are incredibly informative—with U.S. imports by country, legislated rates, and collected duties—but their absence between 1946 and 1964 renders them of little help in analyzing major changes in tariff rates arising from GATT negotiations. Less detailed monthly reports were produced during this period.

Appendix M, but give some highlights of the dataset here.

Our starting point was the tariff scheduled legislated in the Smoot-Hawley Act of 1930. Those tariff rates were also accompanied by trade flow data, which allows us to appropriately address the role of specific tariffs. Between 1930 and 1946 (the year of the first round of the GATT), the U.S. engaged in a slew of bilateral trade negotiations. All changes in tariff rates over this period were recorded in a document produced by the TC in 1946, in anticipation of the first GATT. Each rate change over this period contained a reference to the relevant trading partner, allowing us to systematically document the countries for which negotiations induced regressivity.

We then collected legislated tariffs after each of the seven GATT round concluding before digitized tariff schedules become available in 1989.<sup>25</sup> Over this period, the bulk of changes in tariff rates were made surrounding GATT negotiations, so measuring changes at this frequency should effectively allow us to make statements about when various changes occurred. In Appendix M, we show that the legislated rates we have digitized match other published time series of effective rates of duty at the aggregate level fairly closely.

Figure 2.6 shows snippets of pages of the digitized tariff schedules, for tariffs on *manganese compounds*. All of the data contained on these pages, and pages like them, were manually entered by a data-processing company.<sup>26</sup> In total, we digitized about 2,500 pages of legislated tariffs like the ones shown here. We parsed the rates of duty into specific and ad valorem components using natural language processing (NLP) techniques. For varieties whose verbal description contains a description of its unit value, we parse the descriptions using NLP techniques in order to separate high- and low-value varieties. This, alongside the import data present in the 1930 schedule, are used extensively in section 2.5.2. In section 2.5.3, we also make use of the “column 2” rate of duty which, for the most part, reflects each variety’s 1930 rate of duty (in the case of manganese compounds, 25 percent). Finally, we record the level of indentation of each variety, which is useful in locating sub-varieties (as discussed in section 2.2.1).

---

<sup>25</sup>These occurred in 1949 (Ancey), 1950 (Torquay), 1956 (Geneva), 1960 (Dillon), 1964, (Kennedy), and 1973 (Tokyo). The Uruguay round commenced in 1986, but did not end until after digitized tariff schedules are available.

<sup>26</sup>The company is called Hi-Tech iSolutions LLP.

Figure 2.6: Examples of Digitized Tariff Schedules: Manganese Compounds

1930: Foreign Commerce and Navigation of the U.S.

636 No. 9, PART 2.—IMPORTS ENTERED FOR CONSUMPTION—JUNE 18—DEC. 31, 1930

Class No.	Article	Unit of quantity	Rate of duty	Quantity	Value	Duty	Equivalent ad valorem
GROUP 8.—CHEMICALS AND RELATED PRODUCTS—Continued.							
INDUSTRIAL CHEMICALS—Continued.							
Other industrial chemicals, n. e. s.—Continued.							
Manganese compounds—							
838. 780	Borate.....	Pound.....	25%.....	8, 214	Dollars 638	Dollars 159. 50	Per cent .....
838. 800	Resinate.....	Pound.....	25%.....	22, 110	909	227. 25	.....
838. 810	Sulphate.....	Pound.....	25%.....	734	106	26. 50	.....
838. 820	Manganese compounds and salts, n. s. p. f.....	Pound.....	25%.....	80, 952	3, 649	912. 25	.....

1950: United States Import Duties

UNITED STATES IMPORT DUTIES (1950)

Par. No.	Description	1930 rate	Modified rate	Reference
50	<b>Manganese:</b>			
	Borate, and resinate .....	25% ad val.	20% ad val.	TD 51802
	Sulphate .....	25% ad val.	12½% ad val.	TD 51802
	Other manganese compounds and salts .....	25% ad val.	20% ad val.	TD 51802

1968: Tariff Schedules of the U.S.

Page 234

SCHEDULE 4. - CHEMICALS AND RELATED PRODUCTS  
Part 2. - Chemical Elements, Inorganic and Organic Compounds, and Mixtures

4 - 2 - C  
419.24 - 420.40

Item	Stat. Suffix	Articles	Units of Quantity	Rates of Duty	
				1	2
419.40	00	Manganese compounds:			
419.42	00	Borate.....	Lb.....	9% ad val.	25% ad val.
419.44	00	Sulfate.....	Lb.....	9% ad val.	25% ad val.
		Other.....	Lb.....	12.5% ad val.	25% ad val.

1980: History of the Tariff Schedules of the U.S.

HISTORY OF THE TARIFF SCHEDULES OF THE UNITED STATES ANNOTATED

Notes p. 10  
Schedule 4  
Part 2

Staged Rates

Modifications of column 1 rates of duty by Pres. Proc. 4707 (Tokyo Round), December 11, 1979, 44 F.R. 72348, as amended and Pres. Proc. 4768 (Customs Valuation Code), June 28, 1980, 45 F.R. 45135, as amended (see Preface):

TSUS Item	Prior rate	Rates of duty <u>1</u> /, effective with respect to articles entered, or withdrawn from warehouse, for consumption on and after --							
		January 1, 1980 except as noted	January 1, 1981	January 1, 1982	January 1, 1983	January 1, 1984	January 1, 1985	January 1, 1986	January 1, 1987
419.40	5% ad val.	4.8%	4.7%	4.5%	4.4%	4.2%	4%	3.9%	3.7%
419.42	5% ad val.	4.8%	4.7%	4.5%	4.4%	4.2%	4%	3.9%	3.7%
419.44	7% ad val.	6.7%	6.4%	6.1%	5.9%	5.6%	5.3%	5%	4.7%

Additionally, we digitized just under 400 additional pages of concordances within and across classification systems, which aided in our tracing of tariff rates of specific items back to, or forward from, 1930. These concordances are complete and detailed from 1963 forward (i.e., at the tariff-line level). For the period before 1963, the concordances are a bit cruder, requiring some manual concordance.

## 2.5 Origins of Regressivity: Systematic Analysis

In this section, we broaden the scope of the case study analysis using our newly digitized data in order to show robust and systematic evidence of our main finding: regressivity began in the 1930s and 40s in response to a much different economic landscape and set of policy priorities, and was never fully corrected. We begin, first, in Section 2.5.1 by showing that the time series patterns shown in the case studies hold for a broader panel of representative consumer goods. Using newly digitized data on bilateral trade flows, we highlight the hysteretic nature of rates on these goods—trade patterns have changed substantially since the period in which regressivity was introduced, but the relative rates between goods still reflects policy priorities of the 1930s and 40s.

In Section 2.5.2, we study a panel of goods whose rates were modified between 1930 and 1946, and reach two conclusions. First, the Smoot-Hawley Act rates did not exhibit the regressive pattern in a meaningful way. Second, the bilateral negotiations between 1930 and 1946 (particularly an agreement with the U.K.) introduced a substantial amount of regressivity. We do this using a panel of goods whose varieties are *defined* by their unit value, which is the most direct way to study tariffs on low- and high-value varieties. Recalling that regressivity arises from within-good variation in rates, in Section 2.5.3, we turn to a cross-sectional approach that uses full tariff schedules, and show that within-good variation in rates spiked in the 1930s and 40s, and never fully returned to its 1930 levels.

Figure 2.7: Tariff Rates Over Time for 15 Goods

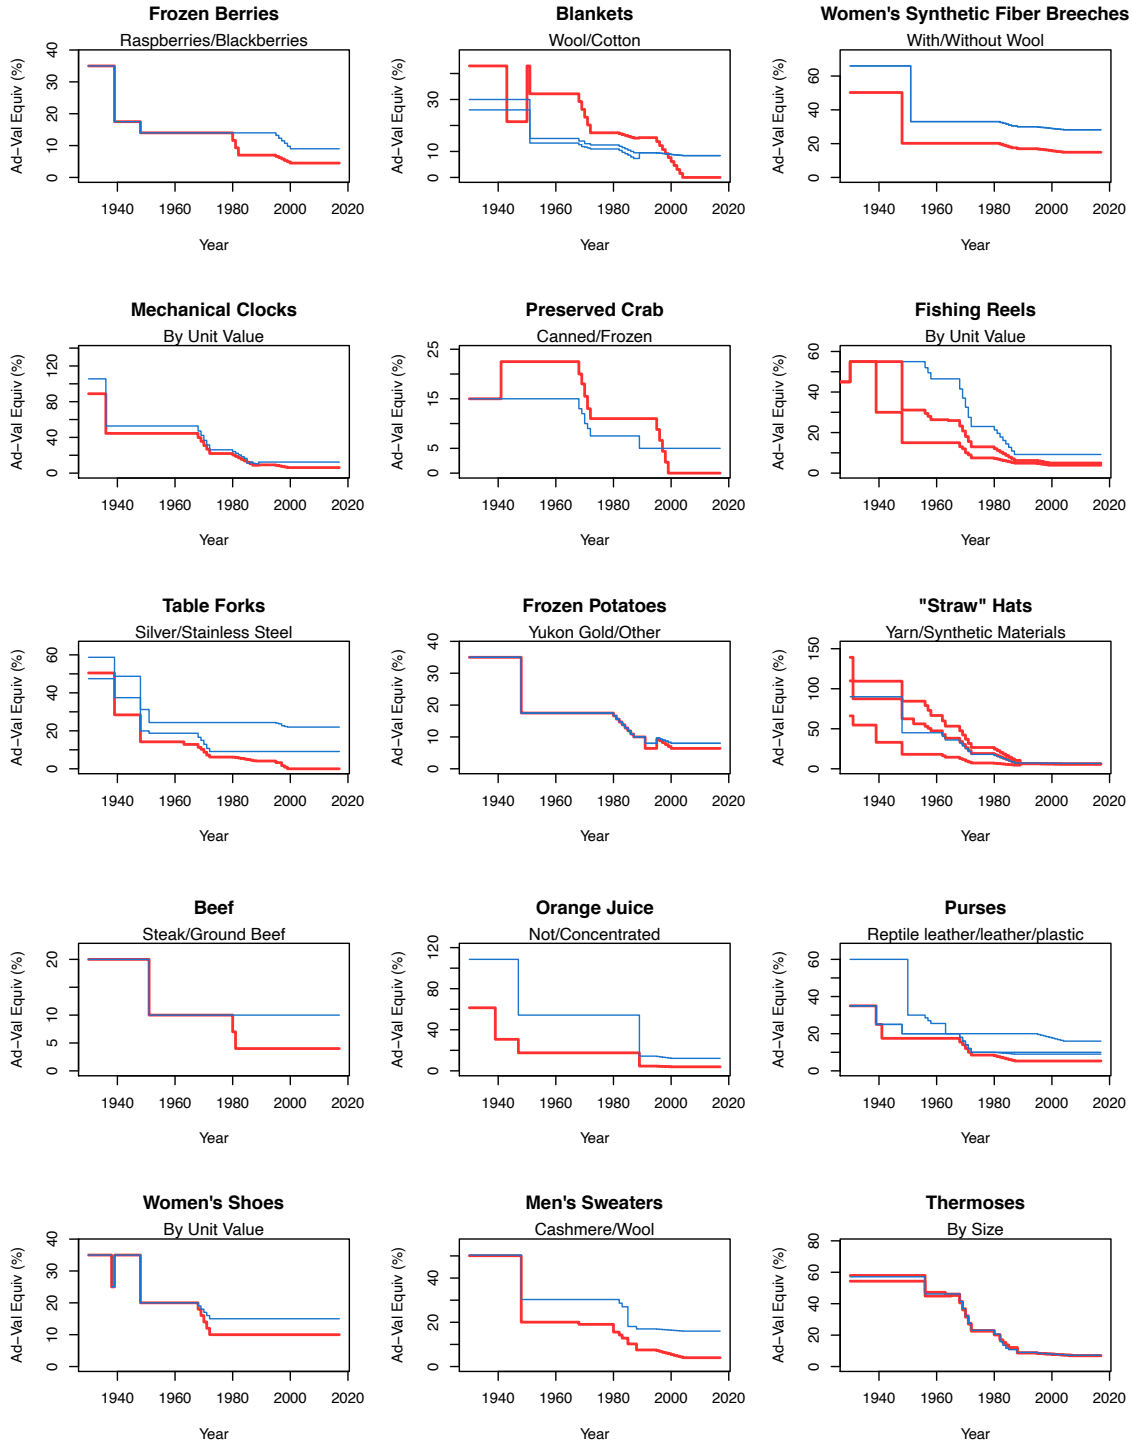



Figure 2.8: Average AVE Tariff Relative to 1930 for 15 Goods



### 2.5.1 Hysteresis in A Long Panel of Consumer Goods

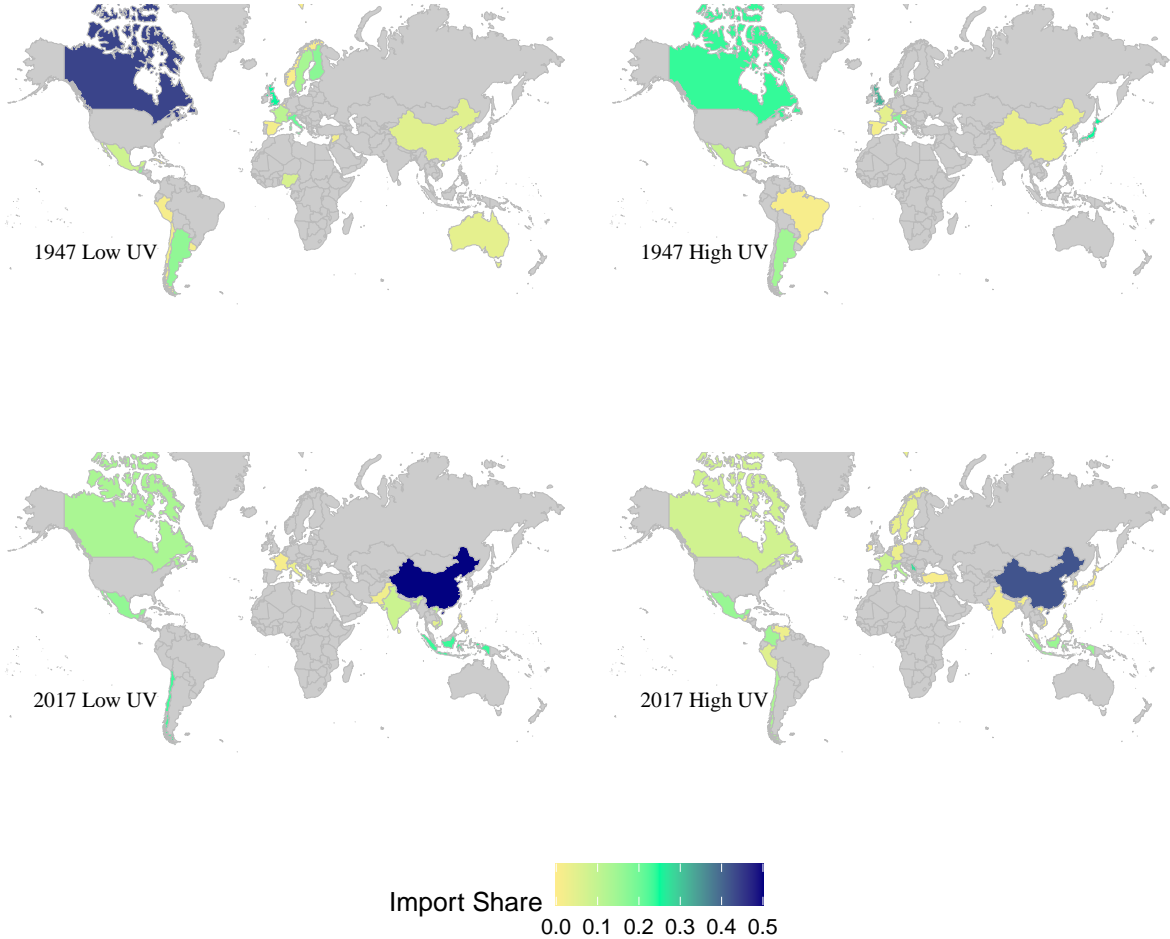
In order to broaden the scope of our historical analysis from fishing reels and forks, we use our newly digitized data to trace variety-level tariff rates back to 1930 for 13 additional goods (and their varieties). We select the goods of focus using criteria based on the approach taken by Crucini (1994), in order to get a representative sample of relevant goods across the U.S. tariff code. Specifically, we begin by selecting the 27 chapters of the HTS that are comprised of at least 75 percent consumer goods by value.<sup>27</sup> We eliminate three of these chapters (HS 46, 66, and 92) because total imports in these chapters failed to surpass 0.1 percent of total imports in 2017. We eliminated five more chapters (49, 97, 83, 57, and 89), because they either contained no dutiable imports or they contained no regressive goods. In the remaining chapters, we selected the top regressive goods, where goods are ranked based on total imports in 2017. In a handful of cases, the top ranked good was not a good candidate for our analysis because it was not composed of obviously substitutable varieties, so we selected the next best good in terms of its import ranking.

The time series of tariff rates for each variety within these 15 goods are plotted in Figure 2.7. While each good and variety has its own story, two general patterns emerge: first, divergences in

<sup>27</sup>As before, the “consumer good” designation relies on the BEC classifications.

tariff rates between low- and high-value goods tend to occur between 1930 and 1980; second, tariff rates on most varieties are relatively stable after around 1990. These two patterns are illuminated in Figure 2.8, which shows the average change in tariff rates relative to 1930 levels on the low- versus high-value varieties of our 15 goods. The divergence between average rates on the two types of varieties begins around 1938—the Anglo-American Trade Agreement of 1938—and the divergence persists over time. At each major round of the GATT tariff rates on all products are reduced, but regressivity is never eliminated. Furthermore, after the Uruguay round of the GATT ended in 1993, average tariff rates are remarkably stable.

Figure 2.9: Source Countries for 15 Goods



One implication of the pattern we find—regressivity emerging in the mid-1900s due to strategic

negotiations—is that unless tariff rates are constantly revisited *at the variety level*, tariff rates that may once have been set optimally to appease a trading partner may no longer be optimal if sourcing patterns have shifted over time. In other words, today’s tariff rates (or, more precisely, the relative tariff rates among varieties of the same good) are functions of the economic landscape of the past, not the present. Returning to the 15 goods, we can see this shifting economic landscape quite clearly. The maps in Figure 2.9 report the average share of imports for the 15 case study goods coming from each country. The top panels show these trade patterns in 1947—a year in the midst of when regressivity emerged for many of these goods. At the time, there were marked differences in the source countries for high- and low-value goods. Low-value goods came predominantly from Canada, while high-value goods had heavy concentrations coming from the United Kingdom, Italy, and Japan. The bottom two panels show import sources for the same low- and high-value varieties in 2017. Of course the most stark change over the 70 year period is that in the present day, the predominant share of *both* low and high unit value varieties comes from China. Moreover, while there are subtle differences in the import sourcing intensities among countries, the low- and high-value maps look remarkably similar to one another, in contrast to the 1947 versions.

### 2.5.2 The Introduction of Regressivity: 1930–1946

Perhaps the most direct way to see the presence of the regressive pattern is to focus on what we call “valued” goods, or goods whose varieties are *defined* by their unit value. For example, fishing reels are taxed differently based on whether their unit values are above \$8.45 each, below \$2.70, or in between. This practice appears for the first time in the tariff of 1816,<sup>28</sup> and provides a set of goods with defined low- and high-value varieties that allows us to circumvent some of the complications in our analysis described in section 2.2. Within this subset of goods, there is no need to construct the definition of a “good,” nor to use unit values to determine which varieties are low- or high-value. We begin our systematic historical analysis by focusing on these valued goods in the Smoot-Hawley Tariff Act. At that time, during the second half of 1930, valued goods represented

---

<sup>28</sup>See the tariffs on cotton products in the “Fourth” paragraph of the 1816 *Act to regulate the duties on imports and tonnage*.

4% of all dutiable goods imported into the United States.

Our analysis of valued goods reveals two important facts about tariffs prior to the GATT (which began in 1947). First, we show that tariffs under Smoot-Hawley were not regressive (at least among this subset of goods), though they appear so due to the profuse use of specific tariffs. Second, we show that the slew of bilateral U.S. tariff negotiations that occurred between 1930 and 1946 introduced an abundance of regressive changes, and much of the within good variation in tariff rates that exists today can be traced back to agreements made during that period.

While the Smoot-Hawley Act ushered in large increases in tariffs on many imported varieties, our analysis of the rates set on valued goods in the Act suggests little evidence for the type of regressive pattern that we document above. There are 59 valued goods contained in the Act, whose legislated tariffs we parsed into their ad-valorem and specific components, as described in section 2.4.<sup>29</sup> Among these 59 goods, 55 contained specific tariffs, and 51 contained both specific and ad-valorem components.

Figure 2.10 shows average tariff rates on the low- and high-value varieties of value-defined goods. The bars labeled “ad-valorem equivalent” show that, relative to the average effective tariff rate for each goods, low-value varieties had a 5% higher effective tariff, and high-value varieties had a 2% lower effective tariff. This finding suggests a pattern of regressivity, but indeed masks an important feature of tariffs at this time: specific tariffs were pervasive. Why? Consider the case of belt buckles, one of the 59 valued goods that became a “valued good” in 1897, largely in response to calls like that from the Alma Button Company of Baltimore in the tariff hearings ahead of the Committee on Ways and Means in 1896/7:

If prices could be exactly determined nothing would seem to be fairer than an ad valorem duty. But unfortunately prices are very much matters of opinion, in which honest men may differ much and rogues much more. Inasmuch as the duty depends on the price, a cheat on the price is a cheat on the duty... Ad valorem duties are mere

---

<sup>29</sup>To define valued goods, we start by parsing the descriptions of all varieties and keep those with the terms “valued” or “Valued” in the description. To group these varieties into “goods,” we keep varieties that are listed sequentially and are at the same level of indentation.

inducements for fraud, and fail to attain the end desired. What is needed is a graduated specific duty.

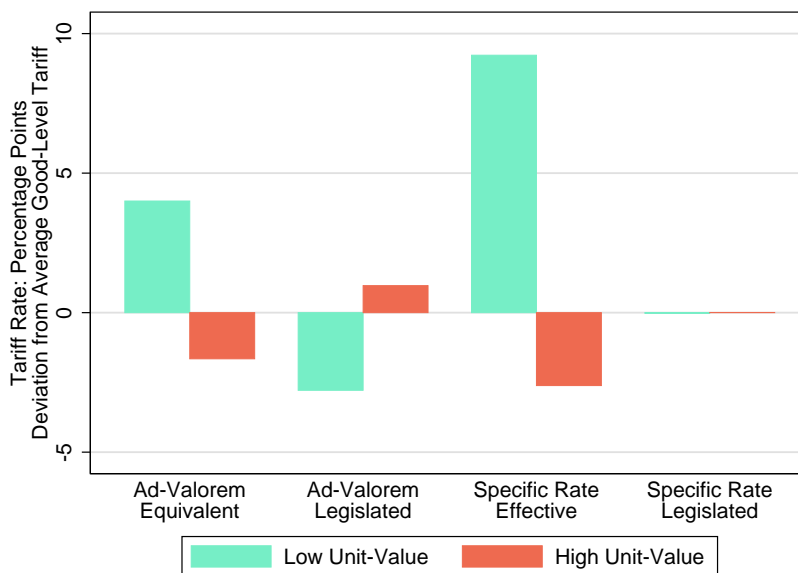
These hearings ultimately culminated in a flat 15% ad-valorem rate on each buckle, and specific tariffs of 5, 10, and 15 cents per 100 buckles on buckles valued below 15¢, between 15¢ and 50¢, and above 50¢, respectively. This graduated schedule of rates was essentially a way to discourage the practice of under-reporting the value of imports for the purpose of circumventing tariffs: the effective ad-valorem rate here is decreasing in price, converging to 15%.

Returning to Figure 2.10, the bars labeled “ad-valorem legislated” show that the ad valorem component of rates on these valued goods do not exhibit the regressive pattern. Instead, rates on low-value goods tend to be lower than rates on high-value goods, and are thus not the source of the regressive appearance of the first set of bars. Instead, the bars labeled “specific rate effective” are to blame. Those bars average the legislated specific rate as a percent of the average imported unit value of each variety. Like in the belt-buckle example, these rates are meant to dissuade under-reporting of prices and provide a minimum amount of tariff revenue per import object. To see this, in the final bars on the right of 2.10, we average the legislated specific tariff as a percent of the legislated unit-value cutoff. At the cutoffs, specific tariffs are (on average) uniform across varieties.

Between the Smoot-Hawley Act and the first round of GATT negotiations in Geneva, the U.S. took part in several bilateral trade negotiations. As discussed in detail by Irwin (2017), many of these agreements were at the urging of former Secretary of State Cordell Hull, who saw trade agreements and the reduction of tariffs as ways to foster international cooperation and peace. For the remainder of this section, we document that many of the tariff modifications made over this period were regressive in nature, with the majority of these regressive modifications coming from the 1939 agreement with the United Kingdom, described in the case study on fishing reels in section 2.3.1.

For this exercise, we use data on legislated tariff changes that took place over the period from 1930 to 1946 (see section 2.4 for more detail). Within these, we found all valued goods by parsing the descriptions of each variety and searching for the term “valued.” We then manually concorded

Figure 2.10: Tariffs on Goods Defined by Unit Value: Smoot-Hawley Tariff




NOTE. The summary statistics here are based on the lowest- and highest-described unit-value variety within each valued good (i.e., when a good has varieties binned into more than two groups by unit value, we retain the first and last bins only). For each variety, we calculate 4 statistics: the ad-valorem equivalent tariff (total duties/imports), the ad-valorem legislated tariff, the “effective” specific tariff (specific tariff/unit value), and the “legislated” specific rate (specific rate/cutoff, where, for low-valued varieties, “cutoff” is the upper-bound of each variety’s unit value range, and vice versa). We then take the difference between each statistic and the (trade-weighted) average value of that statistic for the corresponding good. The light (green) bars show the trade weighted averages of these difference for all low-valued varieties, and the darker (orange) bars show the averages for high-valued varieties.

these varieties back to their Smoot-Hawley rates and import values/quantities from 1930. In many cases, goods that were previously not “valued goods” became so when new varieties were introduced. Thus, over the 1930-1946 period, the number of valued good increased to 105, from the 59 that existed in 1930. In our analysis here, when a good is split into several bins by unit value of its varieties, we retain only the top and bottom bin.

For each variety, we compute ad-valorem equivalent (AVE) tariffs using 1930 unit values (import quantity divided by import value).<sup>30</sup> We then compare these 1930 rates to new rates that were established between 1930 and 1946. Figure 2.11 presents the data in several formats. The left panel contains a scatter plot of 1930 rates on the x-axis against their new rates established over this period on the y-axis. Unsurprisingly, almost all of the modifications that were made are decreases

<sup>30</sup>That is, we sum the ad-valorem component of each variety’s tariff with the quotient of the specific rate and the unit value.

Figure 2.11: Tariff Changes in Valued Goods during 1930–1946



NOTE. This figure shows different representations of changes in ad-valorem equivalent tariff rates between 1930 and 1946. The left panel shows a scatter plot of 1930 rates against new rates established after 1930. The middle panel shows the trade-weighted changes in these rates for high- and low-valued goods, and the right panel shows the number of reductions on these goods by country. As described in the text, ad valorem equivalents are constructed using 1930 import values and quantities, so all changes shown only arise from changes in legislated tariffs.

in tariffs, and many were substantial, consistent with the findings of Bown and Irwin (2016).<sup>31</sup> The second panel confirms that the cuts were substantial, and reductions mostly occurred among high-value goods. The average decrease for high-value varieties was just under 25 percentage points, while the decrease for low-value goods was about 2 percentage points. The final panel presents simple counts of the number of reductions that were made on high-value and low-value varieties. This confirms what can be seen by visual inspection of the scatter plot: most of the reductions that were made were made on the high-value goods in the trade agreement with the U.K.

### 2.5.3 Within-Good Tariff Variation Post 1946

Putting regressivity aside momentarily, our case studies suggest that general variation in tariff rates *within goods* emerged in the mid-1900s and has persisted over time. We test for this pattern across the entirety of the tariff code using digitized tariff schedules from eight years between 1946 and 1980.<sup>32</sup> Specifically, we decompose the variance of tariff rates across the tariff schedule into two components: within-good variation (across varieties of the same good) and across-good variation:<sup>33</sup>

<sup>31</sup>Bown and Irwin (2016) find that tariff rates had come down substantially from Smoot-Hawley levels *prior* to the beginning of the GATT.

<sup>32</sup>The years we use are 1946, 1950, 1952, 1958, 1963, 1968, 1972, and 1980.

<sup>33</sup>“Goods” here, before 1963, are tariff paragraphs. After 1963, we use 5-digit TSUS codes.

$$\text{var}(\tau_{ig,t}) = \frac{1}{N} \sum_{g \in G} \sum_{i \in g} (\tau_{ig,t} - \bar{\tau}_t)^2 = \underbrace{\frac{1}{N} \sum_{g \in G} \sum_{i \in g} (\tau_{g,t} - \bar{\tau}_t)^2}_{\text{across}} + \underbrace{\frac{1}{N} \sum_{g \in G} \sum_{i \in g} (\tau_{ig,t} - \bar{\tau}_{g,t})^2}_{\text{within}} \quad (2.1)$$

where  $i$  indexes varieties,  $g$  indexes goods, and  $N = \sum_{g \in G} |g|$ . We perform this decomposition for each year  $t$  for which we have digitized schedules, and compare the relative within-good variation share in year  $t$  to the within-good variation share in 1930. To perform the 1930 decomposition, we leverage the fact that the 1930 rates for each variety are reported in each year's tariff schedules as the column 2 rate of duty.<sup>34</sup> This allows us to compute the decompositions in each year relative to 1930, without needing to crosswalk tariff lines between that year and 1930.

The results are shown in Figure 2.12. Between 1930 and 1946, the within-good share of variance nearly triples, consistent with our finding that the variation within goods emerged in these early post-Smoot-Hawley trade negotiations. The within-good share of variation declines between 1952 and 1958, and then settles at around 1.5 times its 1930 levels in 1963, where it remains.


## 2.6 Economic Significance

On top of providing a window into the hysteretic nature of tariff policy, tariff regressivity has direct aggregate and distributional implications. In Section 2.6.1, we consider the aggregate implications of the regressive pattern by computing the loss in tariff revenue that would accrue from eliminating the pattern. We then consider the distributional consequences of the pattern in Section 2.6.2. There, we present a reduced-form exercise using the sufficient statistics approach from Borusyak and Jaravel (2021) and regional trade data to show that regressivity amplifies the pro-poor effects of trade liberalization.

---

<sup>34</sup>The U.S. tariff schedule has two main columns. Column 1 reports the MFN tariff rate, and column 2 reports the rates for those countries that do not have normal trade relations status with the United States, which is also the 1930 rate of duty.

Figure 2.12: Within-Good Tariff Variation Share Relative to 1930

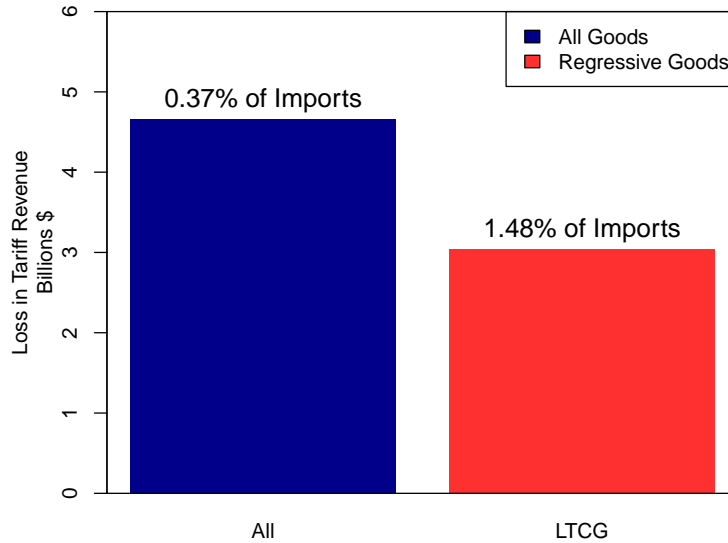


NOTE. This figure shows the within-good share of tariff variation in each year relative to the within-good share of variation in 1930. The variance decomposition is calculated using Equation 2.1

### 2.6.1 Aggregate Implications

As discussed in Section 2.2.3, the economic relevance of the regressive pattern does not hinge on current consumption habits. That said, in this section we ask whether U.S. consumers are meaningfully taxed more heavily on their consumption of low-value goods. To that end, we compare tariff revenues collected with current MFN ad valorem equivalent rates to counterfactual tariff revenues that would be collected if the regressive pattern were eliminated. To construct the counterfactual tariff schedule, within regressive goods we replace the tariff on low-value varieties with the average tariff on high-value varieties. Put differently, assuming that tariffs are fully passed-through into consumer prices and quantities remained unchanged in light of the change in prices, this number represents the amount that consumers would save each year from a reduced tax bill if regressivity is eliminated—an approximation to the loss in consumer surplus from the existence of

Figure 2.13: Revenue Loss from Eliminating Regressive Pattern



NOTE. This figure shows the revenue losses in levels (bars) and relative to total imports within each category (printed numbers) under the counterfactual tariff schedules described in the text.

the pattern. The exercise is an import-weighted representation of Figure 2.2.<sup>35</sup>

Figure 2.13 presents the results. If this change is implemented for all goods, consumers would save nearly \$5 billion per year. If the change is implemented for LTCGs only, the savings would be about \$3 billion, or about 1.5% of imports of LTCGs. This savings is substantial: it represents nearly 20% of total revenue collected on LTCGs. This aggregate figures masks an important aspect of this counterfactual: the gains would likely not be distributed equally, to the extent that lower-income consumers spend relatively larger shares on exactly the more highly taxed varieties. In the next section, we discuss the distributional consequences of the regressive pattern.

## 2.6.2 Sufficient Statistic Calculation

Borusyak and Jaravel (2021) show that import shares in consumer expenditure are sufficient

<sup>35</sup>To see this, suppose that each good  $g$  has a single low- and high-value variety, indexed by  $\ell$  and  $h$ . This graph shows the revenue loss  $\sum_g \sum_{v \in \{\ell, h\}} (\tau_v - \tilde{\tau}_v) m_v$ , where  $\tilde{\tau}_h = \tau_h$  and  $\tilde{\tau}_\ell = \tau_h$  are the counterfactual tariff rates. The revenue loss is therefore  $\sum_g (\tau_\ell - \tau_h) m_\ell$ , or an import-weighted version of the “tariff divergence” measure.

statistics for estimating the welfare effects of trade shocks. Specifically, they show that the welfare effect for an individual with income  $i$  of a reduction in iceberg trade cost on variety  $\omega$ ,  $\partial \log \tau_\omega$ , is given by:

$$\partial \log W_i = \sum_{\omega} s_{\omega}^i \times -\partial \log \tau_{\omega} \quad (2.2)$$

where  $s_{\omega}^i$  is the share of individual  $i$ 's expenditures spent on imported variety  $\omega$ .<sup>36</sup> We estimate equation (2.2) using district-level import data for counterfactual reductions in variety-level tariff rates. The data and calculations are described in more detail below.

### 2.6.3 Data

To estimate equation (2.2), we need data on the expenditure shares of imported variety  $\omega$  in each income quantile  $i$ . We do not directly observe these values, but we can estimate them using regional data on imports and regional data on income and expenditures. For imports, we rely on customs data on variety-level imports by Customs district in 2017. Following Riker (2013), we aggregate districts into 27 import regions such that each region contains at least one customs district and at least one state.<sup>37</sup> District-level import data has its limitations, so we admit that these estimates must be taken with caution. In particular, the import data do not directly identify the location of the final consumer of the imports, but rather the point of entry into the United States from the exporting country. Riker (2013) shows that imports do tend to benefit consumers within the region where they clear customs, so we are cautiously optimistic that these data give us a reasonable proxy for regional import consumption.

We merge the import data at the state level (recall that each import region contains at least one state) with income data from the Current Population Survey's 2017 Annual Social and Economic

---

<sup>36</sup>Borusyak and Jaravel (2021) have a more complex formula that also takes input-output relationships into account, which we abstract from that here.

<sup>37</sup>As in Riker (2013), districts are combined such that regions combine customs districts that have ports in the same state (e.g., Miami and Tampa are in the Florida region), and they also combine states that have ports in the same customs district (e.g., the Maine district includes ports in both Maine and New Hampshire).

Table 2.5: Income and Import Habits

	$\frac{s_{\omega}^i}{\bar{s}}$	$\frac{s_{\omega}^i}{\bar{s}}$	$\frac{s_{\omega}^i}{\bar{s}}$	$\frac{s_{\omega}^i}{\bar{s}}$
$\mathbf{1}\{p_{\omega} > \bar{p}_{g(\omega)}\} \times \ell_i$	-0.0345 (0.0200)	-0.0345 (0.0201)		
$\mathbf{1}\{p_{\omega} > \bar{p}_{g(\omega)}\} \times y_i$			0.192 (0.147)	0.192 (0.148)
Observations	336501	336501	336501	336501
Good FE	No	Yes	No	Yes

NOTE. This table shows the estimates of equation (2.3). Standard errors are clustered at the regional ( $i$ ) level. The columns with “Good FE” include good level ( $g(\omega)$ ) fixed effects. Regional median income,  $y_i$ , is in thousands of dollars, and the low-income share,  $\ell_i$ , is in percent.

Supplement. We then compute the median income of each region. In our analysis below, we will treat each region as being composed of a representative consumer with income given by that region’s median income,  $y_i$ . We normalize each region’s imports so that each region spends 20% (the average value across regions) of its income on imports. Thus, our identification of welfare effects comes from the allocation of spending across imported varieties, rather than between imported and non-imported varieties. To have a sense of whether this treatment of the data provides a sensible approximation of expenditure habits by income, we estimate

$$s_{\omega}^i = \beta [\ell_i \times \mathbf{1}\{p_{\omega} > \bar{p}_{g(\omega)}\}] + \delta_{g(\omega)} + e_{i,\omega} \quad (2.3)$$

where  $s_{\omega}^i$  is region  $i$ ’s imports of variety  $\omega$  divided by total income in region  $i$ ,  $\ell_i$  is the percent of the region’s population that makes below \$20,000 per year,  $\mathbf{1}\{p_{\omega} > \bar{p}_{g(\omega)}\}$  is an indicator for whether variety  $\omega$  has a unit value above its good’s median unit value, and  $\delta_{g(\omega)}$  is a fixed effect for variety  $\omega$ ’s good. The results are displayed in table 2.5. For interpretability, we divide  $s_{\omega}^i$  by its average value of 0.000009. While not terribly statistically significant, the negative estimate of  $\beta$  indicates that as the low-income share increases, relatively less is spent on high-value varieties. Increasing the low-income share by one percentage point decreases the share of income spent on high-valued varieties by 3% (recall, over a base of 0.000009).

## 2.6.4 Estimation

Given actual (gross) MFN ad-valorem tariff rates,  $\{\tau_\omega\}$ , we consider two counterfactual tariff schedules,  $\{\tau_\omega^c\}$ . Before describing the exact counterfactuals, we discuss the procedure for calculating welfare gains over the income distribution using the regional-level import data and equation (2.2). For each region  $i$ , we estimate  $\partial \log W_i$  using equation (2.2), log changes in tariffs  $\{\log(\tau_\omega^c) - \log(\tau_\omega)\}$ , and region  $i$ 's imports of  $\omega$  divided by total income in  $i$ ,  $s_\omega^i$ . We then fit an exponential curve to the relationship between the welfare effect  $\partial \log W_i$  and median income  $y_i$  in each region, and report the fitted curve.<sup>38</sup>

The first counterfactual we consider is an across-the-board reduction in tariffs by 5 percentage points. We call this counterfactual the “uniform reduction” counterfactual and index it by U. The associated tariff schedule has  $\tau_\omega^U = \min\{\tau_\omega - 0.05, 0\}$  for each  $\omega$ . This counterfactual reduces collected duties by 4.4%. The second counterfactual, which we call the “eliminate regresivity” and index by ER, asks how the distribution of welfare gains would change if, in addition to the uniform reduction, tariffs on low-value varieties,  $\omega$ , that fit the regressive pattern were set to the average tariff on the high-value varieties of the same good (within  $g(\omega)$ ). The counterfactual tariff rate is computed as:

$$\tau_\omega^{\text{ER}} = \begin{cases} \min\{\bar{\tau}_g(\omega) - 0.05, 0\} & \text{if } p_\omega < \bar{p}_{g(\omega)} \\ \min\{\tau_\omega - 0.05, 0\} & \text{if } p_\omega \geq \bar{p}_{g(\omega)}, \end{cases}$$

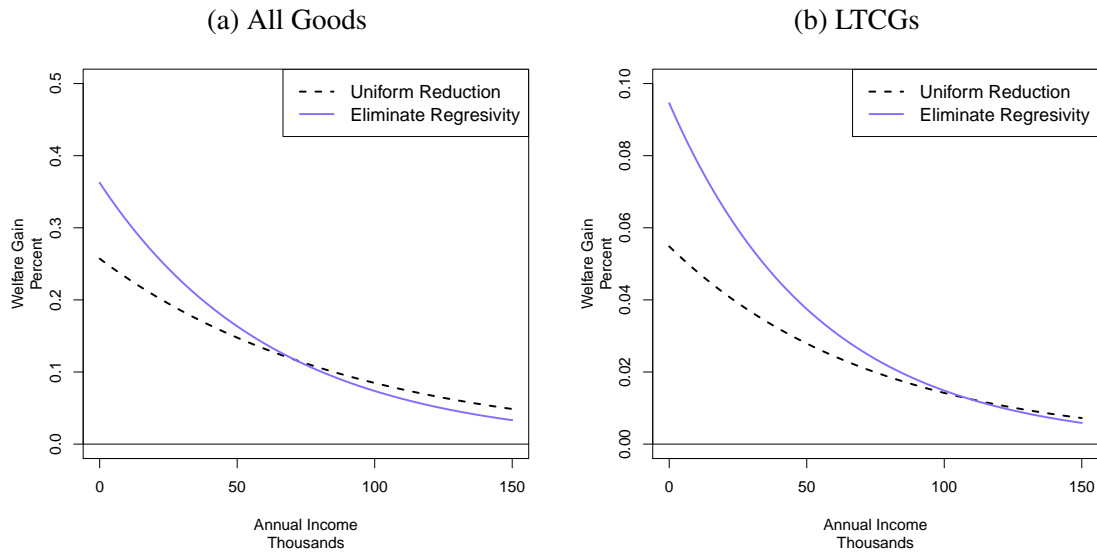
where  $\bar{p}_{g(\omega)}$  is the median unit value of varieties in good  $g(\omega)$ , and  $\bar{\tau}_g(\omega)$  is the average ad valorem equivalent tariff rate on varieties with unit values above  $\bar{p}_{g(\omega)}$ .

The results are shown in Figure 2.14. When performing the counterfactual among all imported varieties, the gain for a consumer with 2017 median U.S. income of \$30,100 is 0.18 percent in the uniform reduction counterfactual, and 0.22% when regresivity is also eliminated. This is at the low end of the range reported by Fajgelbaum and Khandelwal (2016), who estimate that a 5 percent decrease in tariffs on manufacturing imports increases the welfare of the representative consumer

---

<sup>38</sup>An exponential curve fits the data better than a linear one.

Figure 2.14: Welfare Gains under Alternative Counterfactuals



by between 0.2 and 1.3 percent. Among LTCGs, the welfare gain for a uniform 5 percent reduction is 0.036 percent, and this increases to 0.054 percent when we additionally eliminate regressivity. Figure 2.14 shows, however, that welfare gains are not evenly distributed. Consumers at the low end of the distribution see much larger welfare increases than consumers at the top of the income distribution, and this difference becomes especially pronounced in the counterfactual that eliminates regressivity. In other words, the regressive pattern alone has distributional implications.

## 2.7 Conclusion

In this paper, we document a surprising feature of the U.S. tariff code: tariffs are systematically higher on low-value versions of goods relative to their high-value counterparts. We show that this regressive pattern is pervasive across the tariff code, and is not a new phenomenon. Using newly digitized data on legislated tariffs dating back to the 1930 Smoot-Hawley Tariff Act, we show that many regressive goods were born in the 1930s and 40s, and though average tariff rates came down over time, relative rates within goods persisted. Our historical analysis helps to put previous findings on the endogeneity of import classification systems into a dynamic setting: not only is classification endogenous, but classification and therefore tariff rates can be hysteretic as well.

Due to the sheer magnitude and complexity of the tariff schedules, classification and tariff rate decisions are not necessarily revisited or re-optimized in each period, but can persist even as the economic landscape shifts, leaving us with set of tariff rates that reflect policy objectives of a bygone era. Regressive tariffs may have began almost a century ago, but they have modern-day impacts, substantially altering the implied distributional gains from trade.

## Chapter 3: A New Measure of Central Bank Transparency and Implications for the Effectiveness of Monetary Policy

### 3.1 Introduction

Over the years, the world's major central banks have become more transparent in many respects—explicit inflation targets are well-established, policy actions are announced, and forecasts are provided to the public.<sup>1</sup> An aspect of transparency that is more difficult to quantify, however, is the extent to which the publicly-provided rationales for policies reflect the reaction function of the monetary policy committee. As stressed by Woodford (2001, 2005) and numerous speeches by Federal Reserve Chairs,<sup>2</sup> communicating these rationales helps the public to better understand how policymakers might react to different future states of the economy, which in turn gives monetary policy greater control over longer-term interest rates. In this paper, I provide the first high-frequency measure of this “procedural transparency.” With a quantitative measure in hand, we can then tackle the question of whether transparency make monetary policy more effective. The second

---

<sup>1</sup>This chapter is forthcoming in the *International Journal of Central Banking*, and previously circulated under the title “FOMC Responses to Calls for Transparency” (Acosta, 2015). That incarnation of the paper benefited greatly from discussions with my undergraduate honors thesis advisor, John Taylor; the support of my parents and sister: José Acosta, Mary Grebenc, and Lourdes Acosta; and the helpful comments of Lydia Cox, Geoffrey Cox, Marcelo Clerici-Arias, Chris Gust, José Gutierrez, Donna Hunter, Stephanie Kalfayan, Mark Lucianovic, Ellen Meade, the late Allan Meltzer, Elmar Mertens, Kurt von Tish, the Bing Honors College (2013) participants, and seminar participants at the Federal Reserve Board. This version was further improved by participants in the Monetary Economics Colloquium at Columbia, as well as by discussions with Pierpaolo Benigno (editor), Lydia Cox, Juan Herreño, Jennifer La’O, Frederic Mishkin, José Montiel Olea, Emi Nakamura, Andrea Prat, Ricardo Reis, Martin Rotemberg, Joseph Saia, Jón Steinsson, and several anonymous referees. A huge thanks goes to Stephen Hansen and Lucas Husted, who graciously provided me with data, and Blaine Helleloid, who provided outstanding research assistance. This work was supported financially by the Alfred P. Sloan Foundation Pre-Doctoral Fellowship in Behavioral Macroeconomics awarded through the NBER.

<sup>2</sup>See, for example, Bernanke (2010, 2013, and Yellen (2013). The aspects of transparency described by these papers is not *exactly* “procedural transparency” as used here, but it is argued below that procedural transparency is the relevant form of transparency for explaining the rationale of policy, which is what these papers and speeches broadly consider.

part of this paper provides an answer in the affirmative.

I construct the measure of procedural transparency—henceforth referred to as “transparency” unless otherwise noted—using the documents that Federal Open Market Committee (FOMC) releases to the public as records of its policymaking meetings. Since 1976, the FOMC has consistently recorded nearly-verbatim records of its meetings in documents called the *transcripts* and shorter summaries called the *minutes*. The transcripts, while detailed accounts of how the Committee comes to decisions, are not released until at least five years after a meeting has taken place. (Before 1993, they were not expected to be released at all.) The lag between meetings and the release of the transcripts renders much of the information they contain stale in terms of understanding the Committee’s current thinking. The minutes, on the other hand, are released with a much shorter lag—today, three weeks. My measure of transparency takes advantage of this timing: it is the similarity of the minutes and transcripts of each meeting, computed using natural language processing (NLP) techniques.

Minute-transcript similarity and the distribution of the topics in the minutes and transcripts are most-strongly associated with FOMC leadership. Minute-transcript similarity was at its highest during the late Greenspan years (early 2000s), and increased noticeably after 1993, the year in which the FOMC began publishing the records of its meetings. This is consistent with predictions from earlier studies of this event. The distribution of topics in the transcripts also changed significantly in 1993. Minute transcript similarity is not predictable by macroeconomic variables, though changes in the distributions of the minutes and transcripts are. Higher transparency is also weakly associated with monetary shocks—as measured by several authors—that are smaller in magnitude, lending credence to the notion that the measure helps to inform the public about the Fed’s policy. In sections 3.2 and 3.3, I describe the construction of the measure and its properties, respectively, to establish that it provides a meaningful measure of Fed transparency.

In its most literal interpretation, minute-transcript similarity captures the overlap between the distributions of topics discussed in the transcripts, and those discussed in the minutes. My measure should therefore be interpreted as a way to understand the divergence between what receives the

Committee’s attention when a decision is being made, and how that thinking is described.<sup>3</sup> Given a higher level of transparency, then, the public and financial market participants should be able to better-predict the Fed’s policies, as elevated transparency implies that the Fed’s communications are providing clearer insights into policymakers’ thinking. This motivates the main empirical question of the paper: how does transparency affect the effectiveness of monetary policy, as measured by the pass-through from short-term nominal rates to long-term nominal *and* real rates?

Section 3.4 turns to the role that transparency plays in determining the effectiveness of monetary policy. Specifically, I show that the monetary policy shocks of Nakamura and Steinsson (2018) have larger effects on real interest rates when transparency is elevated. In sections 3.5.1 and 3.5.2, I show that these results are primarily driven by transparency about monetary policy strategies conditional on the economic outlook. Additionally, the effects of monetary policy estimated by Nakamura and Steinsson (2018) are shown to be slightly downwardly biased. This arises because some of the larger monetary shocks have been delivered at times when transparency is low. But, it is precisely when transparency is low that monetary policy shocks have smaller effects on real interest rates, possibly because the public cannot make as much sense of the short-run surprises as they relate to the path of future rates. These findings are robust to concerns that transparency may be proxying for some other variable—forward guidance, the state of the economy, and public uncertainty about monetary policy.

## 3.2 Measuring Procedural Transparency

---

<sup>3</sup>Note that a clear, concise, and informative summary of a long discussion—a discussion possibly filled with tangents, misunderstandings, and other banter—should not be expected to cover the same topics in the same proportion, though my proposed measure would penalize such deviations. Variation in minute-transcript similarity could arise from the noise present in natural language/conversation, variation in how difficult a meeting is to describe (say, for example, because of a complicated policymaking environment), or strategic considerations regarding the transmission of information. That said, to my knowledge, there is no evidence that the minutes are anything but a forthright effort to summarize the transcripts (though, such evidence might be hard to find). Additionally, I find in section 3.5.2 that my results are robust to using measures of transparency about obviously-meaningful topics—i.e., measures that should be less-affected by these concerns.

### 3.2.1 Previous Measurements

The literature on central bank transparency and communication started in earnest at the turn of the century—Blinder et al. (2008) provide a thorough survey of this literature through 2008. This paper contributes to a branch of this literature concerning the measurement of transparency—most recently treated by Dincer and Eichengreen (2014), preceded by Eijffinger and Geraats (2006),<sup>4</sup> who based their measures on the Geraats (2001, 2002) definitions of different aspects of central bank transparency. This paper centers around a particular component of Geraats’ *procedural transparency* that concerns central bank accounts of deliberations, the measurement of which has, thus far, focused primarily on fairly aggregated and slow-moving measures of the timeliness and informativeness of central bank communications.<sup>5</sup> For example, the relevant measure of procedural transparency in Dincer and Eichengreen (2014) is a binary indicator of whether “the central bank give[s] a comprehensive account of policy deliberations (or explanations in case of a single central banker) within a reasonable amount of time[.]” In this paper, the use of natural language processing techniques allows these shortfalls to be circumvented by using the text of each FOMC meeting. Because the measure is constructed from text, it reflects fairly detailed changes in communications and transparency (it is real valued, as opposed to taking a discrete number of values). Because it changes at every FOMC meeting, it gives a high frequency measure of transparency that, in practice, changes much more often than previous measures that primarily capture large regime changes in communications policies.

### 3.2.2 How to Measure Transparency

My proposed measure of procedural transparency is the similarity between the minutes and transcripts of each FOMC meeting, which I will refer to as “minute-transcript similarity” or simply “transparency” for the rest of the paper. The transcripts contain a nearly-verbatim record of each

---

<sup>4</sup>Dincer and Eichengreen (2014) also thoroughly discuss previous literature.

<sup>5</sup>See appendix O for a discussion about the different types of transparency relevant for a central bank. Procedural transparency, as defined by Geraats (2002) is the description of how monetary policy decisions are made, which is achieved in part through the publication of records of the deliberative process.

FOMC meeting, yet are not released until at least five years after an FOMC meeting has taken place—before 1993, they were not released at all.<sup>6</sup> The minutes are shorter summaries—typically 10-20 pages to summarize 100-200 page transcripts—released three weeks after each meeting has taken place.<sup>7</sup> Evidence suggests that the minutes are intended to be accurate portrayals of what was discussed at each meeting, and are not intended to obfuscate the content of the discussion. Chair Yellen, questioned about this in her June 2016 press conference, responded that “[T]he minutes are always—have to be an accurate discussion of what happened at the meeting.”<sup>8</sup> In addition, the FOMC has to vote on the minutes, presumably reducing the possibility of systematic obfuscation.

The first step in computing minute-transcript similarity is to represent each transcript and minutes as a distribution over a finite number of topics, using Latent Dirichlet Allocation. The similarity between the minutes and the transcripts for a particular meeting is the Kullback-Leibler similarity of the distributions of the two documents—a measure that lies in the interval  $[0, 1]$ . The measure bears a striking resemblance to Chair Greenspan’s interpretation of how the public understood FOMC communications, which he voiced in the September 2003 FOMC meeting: “[A] number of those in the market don’t listen to the subtleties; they just take note of how much time we are spending talking about a particular subject” (Transcripts, 1976–2008).

Before detailing the procedure used to compute minute-transcript similarity, it is important to understand what this measure is, and its potential shortcomings. The measure represents the extent

---

<sup>6</sup>“Nearly verbatim” refers to the fact that the exact words are lightly edited. From the Federal Reserve Board’s website: *The most detailed record of FOMC meeting proceedings is the transcript. Beginning with the 1994 meetings, the FOMC Secretariat has produced the transcripts shortly after each meeting from an audio recording of the proceedings, lightly editing the speakers’ original words, where necessary, to facilitate the reader’s understanding. Meeting participants are given an opportunity within the subsequent several weeks to review the transcript for accuracy.*

*For the meetings before 1994, the transcripts were produced from the original, raw transcripts in the FOMC Secretariat’s files. These records have also been lightly edited by the Secretariat to facilitate the reader’s understanding. In addition, where one or more words were missed or garbled in the transcription, the notation “unintelligible” has been inserted. In some instances, words have been added in brackets to complete a speaker’s apparent thought or to correct an obvious transcription error or misstatement.*

<sup>7</sup>The release lag has changed over the years—see N for more details.

<sup>8</sup>This is not a recent sentiment, in 2003, Kansas City Fed President stated in an FOMC meeting that “the minutes are our representation of the discussions that occurred at the meeting” (Transcripts, 1976–2008, Sept. 2003). Several other pieces of anecdotal evidence are available, including discussions I have had with Federal Reserve officials and a former Fed governor.

to which the content of the transcripts is reflected proportionately in the minutes.<sup>9</sup> Because LDA represents each document as a distribution over topics, the mass placed on each topic describes the amount of the document devoted to that topic. Thus, only when the minutes devote exactly the same amount of space to each topic as the transcripts will the measure equal unity. The view taken in this paper is that it is not the job of the minutes writers to editorialize the FOMC’s discussion—e.g. to eliminate the side of a debate that does not ultimately “win”—but instead to convey the discussion accurately. Put differently, minutes that fully communicate FOMC discussions are taken as transparent.

That said, the inclusion of obviously-irrelevant discussions—e.g. “when should we break for lunch”—should not be a necessary condition for minutes to be transparent. Of course some divergence between the two documents should be expected—while a conversation might be centered around a topic, the actual words used or topics discussed might only-noisily represent that topic—owing to, for example, digressions or misunderstandings. The underlying assumption I make is that this noise is fairly constant over time, only affecting the level of my measure and not its changes. In section 3.5.2, I show that my empirical findings are robust to more-narrowly defined measures of transparency that are constructed using economically important topics (and should thus be relatively free from this type of noise). This is a benefit of using the fairly complex language model described in the next section—it allows for documents to be analyzed about detailed topics and thus separate out the types of discussions that can add noise to my measure.

### 3.2.3 Language Model

Several steps are involved in computing the topic distribution for each FOMC transcript and minutes, with the ultimate goal being a representation of a document into well-understood topics in a high-frequency, holistic, and interpretable way. A drawback of the approach used here—along with the vast majority of NLP techniques—is that documents are represented as “bags of words,” i.e., the order of words does not matter. The only features of a document that are retained are

---

<sup>9</sup>Note that the measure does not consider the similarity of documents *at different meetings*, only the similarity of documents related to the same meeting.

counts of the number of times that each unique word appears in that document. This is a necessary evil when working with a large body of text—the 600 documents used in the analysis here contain 13 million individual, and 170,000 unique, words.<sup>10</sup>

After converting the documents into computer readable formats, these raw text files are “pre-processed” using several techniques that are standard in working with natural language. This preprocessing achieves three goals. The first goal is to reduce the effect of errors that might arise from working with a data-source that may contain typographical errors or other errors arising from the fact that several of the documents had to be converted from typewritten documents. To that end, only the letters of the alphabet are retained, and every unique word must appear at least three times over the entire corpus, otherwise it is dropped. Words shorter than three characters or longer than fifteen are also dropped—the intention of the latter being to remove words that may have been accidentally concatenated.

The second goal of the preprocessing is to reduce the noise that arises from grammatical constraints: the words *increase*, *Increase!*, *increasing* and *increased* all convey essentially the same meaning, yet a simple numerical representation of the words in the document might treat them as completely different words, since it knows no better. To that end, words are stemmed to their lexical root, so that in the example above, every occurrence of the “increase” words are stemmed to *increas*.<sup>11</sup> Terms in a “stoplist” are also excluded. As is customary, this list contains common words that contribute little meaning to the documents, since they are used so often. The excluded words are the “generic”, “dates and numbers,” and “geographic” lists from Loughran and McDonald (2011), who carefully constructed these lists to be relevant in a context of finance.

The final goal of preprocessing is to reduce the noise that arises when ideas need to be mapped into words, and vice versa. This goal is addressed via an application of Latent Dirichlet Allocation, developed by Blei et al. (2003).<sup>12</sup> Because LDA enjoys widespread use in the NLP community, and

---

<sup>10</sup>This is based off of simply splitting the documents by whitespace, i.e., with no preprocessing.

<sup>11</sup>Stemming is performed using the Lancaster Stemmer as implemented in Python’s Natural Language Toolkit.

<sup>12</sup>At issue here are the problems of synonymy and polysemy. Polysemy occurs when one word can describe many concepts. For example, polysemy would lead the documents [I read a book.] and [I’ll book a hotel.] to look more

even within the Economics literature, the treatment here is brief. First, the observed corpus contains words, with  $w_{d,n}$  being  $n^{\text{th}}$  word in the  $d^{\text{th}}$  document ( $d \in \{1 \dots, D\}$ ), where each document has  $N_d$  (so that  $n \in \{1, \dots, N_d\}$ ,  $\forall d$ ). This is all that is observed. LDA posits each document as a distribution over a fixed number,  $K$ , of topics— $K$  is chosen by the researcher. Topics are in turn distributions over the  $V$  unique terms in the corpus. More precisely, each document  $\delta_d$  is a draw from a Dirichlet distribution, a distribution over vectors that lie in the  $K$ -simplex. The distribution,  $\delta_d$ , that is drawn from the Dirichlet is a latent variable. The same is true of topics: each topic,  $\phi_k$  is a draw from a  $V$ -dimensional Dirichlet. With a topic distribution in hand, each observed document (which has a fixed length,  $N_d$ ) is populated one word at a time. For the  $n^{\text{th}}$  slot of document  $d$ , a topic is drawn from a multinomial distribution, with parameter  $\delta_d$ . Thus, if  $\delta_d$  is heavily concentrated on topic 1, then several words will be drawn from topic 1. The drawn topic,  $z_{d,n} \in \{1, \dots, K\}$ , is then used to draw a word from a multinomial distribution with parameter  $\phi_{z_{d,n}}$ . So, if a topic has a distribution that places heavy weight on “whale,” then “whale” will come up often when that topic is drawn, and will thus show up often in documents that have a high probability placed on that topic. The implementation here—including the choice of the priors for the Dirichlet distribution—follows very closely that of Hansen et al. (2017), who estimate LDA on a subset of the FOMC transcripts that I consider using a Gibbs-sampler outlined in Griffiths and Steyvers (2004).<sup>13</sup> The number of topics,  $K$ , is set to 50, chosen using a five-fold cross-validation technique similar to that outlined in Hansen et al. (2017) and described in P.

Rather than estimating the topic model over the complete documents in my corpus, I instead begin by splitting these documents (the minutes and transcripts) into sentences using a grammatical sentence parser.<sup>14</sup> I then estimate the word and topic distributions over every sentence in the FOMC

---

similar than an English-speaking human might think. Next, synonymy occurs when any concept can be expressed using many different words. The document [I’ll make lodging arrangements.] and the hotel document from above would look dissimilar, despite conveying the same idea.

<sup>13</sup>And very graciously provided by the authors at <https://github.com/sekhansen/text-mining-tutorial>.

<sup>14</sup>I parse the minutes and transcripts into sentences using the English probabilistic context-free grammar developed by Klein and Manning (2003) as implemented in the Stanford Parser Java package. This sentence parser uses rules of English grammar to split sentences, as opposed to simple rules based solely on punctuation. This, for example, avoids erroneously splitting sentences at decimal points or after abbreviations(e.g. “Ms.”).

Table 3.1: Topic Descriptions

Topic	Top 10 Words										Dissent	FFR	GDP	Inflation
1	level	unemploy	percent	rate	employ	low	averag	high	declin	remain				
2	credit	bank	spread	loan	debt	investor	mortgag	bond	larg	money			✓	
3	hous	sale	sector	inventori	activ	declin	construct	home	start	weak			✓	
4	meet	statement	altern	committe	chang	languag	fomc	paragraph	direct	word	✓		✓	
5	think	make	commun	public	view	way	help	decis	use	forward	✓			
6	market	financi	labor	particip	improv	condit	forc	gener	substanti	note	✓			
7	purchas	asset	treasuri	end	secur	program	hold	back	reduc	size			✓	
8	spend	busi	consum	invest	incom	capit	household	confid	increas	continu	✓			
9	risk	balanc	concern	view	side	possibl	sheet	downsid	outlook	signific				
10	time	take	need	think	begin	adjust	process	place	step	normal				
11	chang	littl	move	think	bit	probabl	reason	case	quit	fact		✓	✓	
12	like	look	see	come	go	time	forward	back	number	reason				
13	polici	monetari	action	fiscal	accommod	tighten	eas	appropri	addit	support	✓	✓	✓	
14	line	dollar	right	panel	lrb	rrb	shown	foreign	show	left		✓		
15	term	long	expect	run	longer	short	consist	near	time	object	✓		✓	
16	growth	continu	economi	pace	slow	moder	product	expans	trend	econom		✓	✓	
17	reserv	bank	oper	feder	system	account	open	central	facil	direct	✓		✓	
18	recent	data	month	suggest	year	past	indic	survey	seen	evid				
19	peopl	lot	get	work	think	way	talk	just	say	thing	✓			
20	price	increas	cost	higher	pressur	rise	energi	oil	wage	declin		✓	✓	✓
21	economi	uncertainti	problem	situat	world	face	shock	deal	event	creat				
22	forecast	project	growth	real	gdp	greenbook	staff	point	half	revis	✓			
23	effect	demand	import	factor	reflect	part	posit	suppli	export	state		✓	✓	
24	inflat	expect	percent	measur	core	remain	pce	rang	higher	low				✓
25	rate	fund	point	interest	basi	feder	target	lower	fed	rais	✓			
26	report	good	product	industri	firm	district	manufactur	order	contact	nation	✓			
27	differ	use	model	base	output	gap	rule	structur	estim	actual				
28	question	presid	issu	chairman	discuss	comment	said	governor	mention	ye			✓	
29	think	go	want	say	know	just	dont	get	realli	way			✓	
30	econom	period	committe	outlook	member	develop	stabil	inform	intermeet	consider				✓

minutes and transcripts between 1976 and 2014. In the notation established above, then,  $d$  indexes sentences so that each sentence has an estimated distribution over topics. With these estimated distributions, I then estimate the topic distributions for each transcript as a whole ( $\theta_t \in \mathbb{R}^K$ ), and each of the minutes as a whole ( $\mu_t \in \mathbb{R}^K$ ). This follows the approach of Hansen et al. (2017), the purpose being to have each document focused on a small number of topics, in hopes that the latent topics can be determined more easily.<sup>15</sup>

The “top 10 words” column of 3.1 contains the ten words that most-prominently contribute to the composition of each topic. Formally, these are the words corresponding to the ten largest elements of  $\phi_k$  for each topic  $k$ . For purposes of interpretation I show the estimated topics for a smaller model with  $K = 30$  (my empirical results in section 3.4 are nearly unchanged quantitatively using the larger model) that I estimated over the sample beginning in 1995 since that forms the basis of my empirical estimates in section 3.4. In general the topics seem intuitive—the first four

<sup>15</sup>This substantially increases  $D$  to about 260,000, though there are 300 minutes and 300 transcripts. This approach of using already-estimated topics to estimate the topic distribution of an excluded document underlies the right panel of figure O.1, with the excluded document reading “transcripts minutes record policy actions memorandum discussion communications.”

topics might be subjectively called topics about labor, credit markets, housing markets, and policy statement language. Not every topic is directly related to an economic concept—the fifth topic contains words that might be used in a debate regarding policy communications.

To better aid in categorizing the estimated topics, the last four columns of table 3.1 indicate which of the topics are useful for “predicting” a few external observable variables, denoted by  $e_t$ . Specifically, I gather the number of dissenting votes at each FOMC meeting from Thornton and Wheelock (2014); and the change in the target of the Federal Funds rate, GDP growth, and inflation from FRED.<sup>16</sup> I then estimate which topics are useful predictors of these variables by estimating the following regression using the LASSO objective function

$$e_t = \beta_0 \mathbf{1}\{\text{ELB}\}_t + \sum_{k=1}^K \beta_k \delta_t^k + \text{error}_t$$

where  $\delta_t^k$  is the estimated presence of each topic  $k$  in document  $d$ . I estimate this regression twice: once using the presence of topics in the minutes (i.e. replacing  $\delta_t^k$  with  $\mu_t^k$ ) and once using their presence in the transcripts (i.e. replacing  $\delta_t^k$  with  $\theta_t^k$ ). I take the union of the selected topics between these two regressions. I select the LASSO regularization parameter using ten-fold cross validation.<sup>17</sup> This again helps to shed light on the estimated topics. The selected topics for inflation are the most intuitive, with words like “price, inflat, econom” being the top words for the selected topics (20, 24, and 30). Instead topics about appropriate communication and policy (4, 5, 25), and the general outlook for the economy (most of the other selected topics) can predict the number of dissents at each meeting.

---

<sup>16</sup>In terms of FRED mnemonics, the FFR target is DFEDTAR when it is available and the midpoint of DFEDTARL and DFEDTARU when it is not; GDP growth is the four-quarter difference in the log of GDPC1; and inflation is the twelve-month difference in the log of PCEPI.

<sup>17</sup>I include a time-varying constant since the FFR is one of my target variables. I implement the estimation and cross-validation using the LassoCV module in Python’s sklearn package; the sample is split into ten disjoint subsets that are the same for each  $e_t$ .

### 3.2.4 Transparency Index Definition

Ultimately, the object of interest in this paper is not the topics themselves, but rather the relative entropy of the minutes for the transcripts (i.e., the Kullback-Leibler similarity) and the entropies of the minutes and transcripts on their own. Specifically, given the topic distribution of the transcripts for the FOMC meeting occurring at time  $t$ ,  $\theta_t$ , and the topic distribution of the minutes for the same meeting,  $\mu_t$ , define the three quantities:

$$\begin{aligned} \text{Minute-Transcript Similarity:} \quad & \tau_t \equiv \exp \left[ - \sum_{k=1}^K \mu_t^k \ln \left( \frac{\theta_t^k}{\mu_t^k} \right) \right] \\ \text{Entropy of the Minutes:} \quad & H(\mu_t) \equiv - \sum_{k=1}^K \mu_t^k \ln \left( \mu_t^k \right) \\ \text{Entropy of the Transcripts:} \quad & H(\theta_t) \equiv - \sum_{k=1}^K \theta_t^k \ln \left( \theta_t^k \right) \end{aligned}$$


with  $\theta_t^k$  being the  $k^{\text{th}}$  element of  $\theta_t$ , and analogously for  $\mu_t$ . The first measure is the negative exponential of the Kullback-Leibler divergence, a distance function for distributions—intuitively it describes the information loss from assuming the truth is the minutes when it is really the transcripts. The other two measures, the entropies of the minutes and transcripts, are the expected values of the information content of a random variable that is distributed according to  $\mu_t$  and  $\theta_t$ . Entropy in this case achieves a maximum when both distributions place equal mass on each topic, i.e.,  $\mu_t^k = \frac{1}{K}, \forall k$ , and decreases as mass moves away from certain topics and concentrates on others.<sup>18</sup> Thus, entropy in this context can be cast intuitively as a measure of how dispersed a conversation/document is—the lower the entropy, the more concentrated the discussion.

For the construction of the procedural transparency measure, the body of documents under consideration consists of all Records of Policy Actions (ROPA, an older version of the minutes), and minutes and transcripts from meetings physically held in Washington, D.C., between April 1976

---

<sup>18</sup>Information about a variable drawn from a uniform distribution is more valuable than information about a variable drawn from a point mass—one already know where the point mass is, but has no idea where the uniform variable is.

Figure 3.1: Procedural Transparency



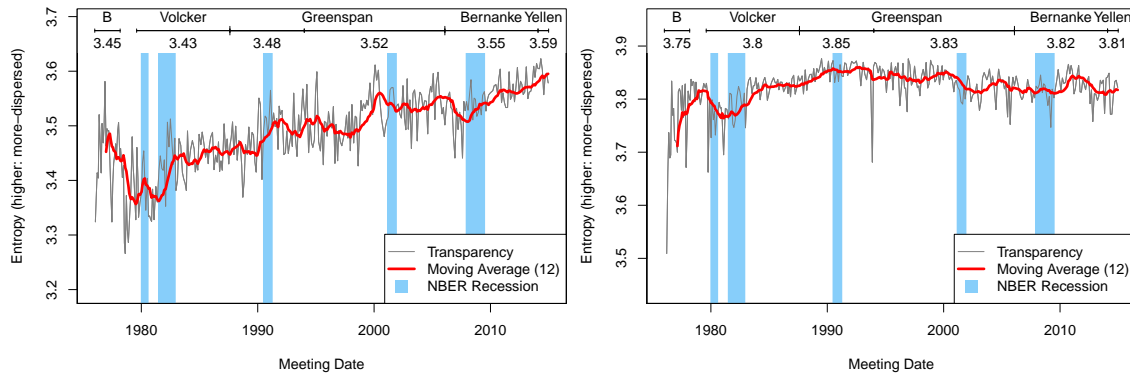
and December 2014, where procedural information (voting records, attendance) is removed.<sup>19</sup>

### 3.3 Notable Features of the Minutes, Transcripts, and Procedural Transparency

This section presents my measure of procedural transparency and three exercises meant to better understand it. First, in section 3.3.1 I compare the measure to several external variables, in order to understand its systematic components. My primary finding is that fluctuations in the measure, while correlated with some of these external variables, is not a proxy for something simpler. Next, in section 3.3.2 I highlight, anecdotally, that the measure captures meaningful differences between the minutes and transcripts. Finally, in section 3.3.3 I argue that my measure—though not directly observable in real-time—might be roughly observable to the public. This is done using a newspaper-based measure of central bank transparency.

<sup>19</sup>See N for a discussion of the various documents released by the FOMC since its inception. For the modern-day minutes (1993–present), all words prior to the paragraph that typically begins with “The information reviewed at the  $x$  meeting...” (now labeled “Staff Review of the Economic Situation”) is removed in order to make these documents look like the ROPA. This also keeps the content of the minutes looking relatively similar over the year, since the first meeting of each year contains discussions of procedural matters (see Meade et al. (2015) for more information on the content of the minutes).

Figure 3.2: Entropy of the Minutes (Left) and Transcripts (Right)



### 3.3.1 Correlations with External Variables

Figure 3.1 presents the time series of minute-transcript similarity, and figure 3.2 contains the entropies of the minutes and transcripts. All three measures are rather noisily distributed around slower moving trends, shown in the images as twelve-meeting trailing moving averages. Owing to concerns—mentioned in section 3.2.2—that the raw measures are likely influenced by idiosyncratic noise arising from the noise inherent in natural language, the moving averages of these series are the main measures considered for the rest of the paper.

In order to analyze the series more carefully, tables 3.2 and 3.3 present the results of regressing the standardized measures on several variables. The first, table 3.2, shows the regression of these communications variables on the other variables one at a time. The lessons here are largely consistent with the results in table 3.3, which estimates the coefficients jointly for variables that span the entire 1976–2014 sample. In table 3.2 the variables are standardized so that the coefficients represent correlation coefficients; in table 3.3, measures are scaled so that they can be interpreted as the number of standard deviations by which the measure moves when the variable in the row increases by one unit.<sup>20</sup> The regressions highlight some notable features of the series.

First, the three measures all contain positive linear trends to different degrees of statistical significance—between the two tables, the minutes and transcripts have generally increased in their breadth of coverage (higher entropy), and transparency has generally increased. One no-

<sup>20</sup>Over the sample period used for that regression.

Table 3.2: Regression Coefficients for Communications Variables

	Transparency		Minutes Entropy		Transcripts Entropy	
		MA		MA		MA
$t$	<b>2.50</b>	<b>3.24</b>	<b>2.63</b>	<b>3.23</b>	1.07	1.41
Transparency MA	<b>0.33</b>	<b>0.43</b>	0.09	0.14	0.18	0.19
Minutes Ent. MA	0.13	0.17	0.26	0.43	0.17	0.16
Transcript Ent. MA	0.04	0.09	0.08	0.07	<b>0.61</b>	<b>0.92</b>
TT State	-0.14	-0.11	-0.10	-0.07	0.12	0.30
RR Shocks $_{t+1}$	-0.03	-0.04	-0.03	-0.02	0.07	0.06
RR Shocks $_{t+1}$	-0.08	<b>-0.16</b>	<b>-0.13</b>	<b>-0.22</b>	<b>-0.13</b>	<b>-0.21</b>
NS Policy Shock $_{t+1}$	0.06	0.13	0.06	0.11	0.06	<b>0.22</b>
NS Policy Shock $_{t+1}$	0.13	-0.01	-0.05	<b>-0.32</b>	0.08	-0.03
FFR	<b>-0.55</b>	<b>-0.77</b>	<b>-0.70</b>	<b>-0.88</b>	-0.20	<b>-0.37</b>
Change in FFR	-0.07	-0.04	-0.09	-0.04	0.01	-0.04
Change in FFR	<b>-0.17</b>	<b>-0.23</b>	<b>-0.27</b>	<b>-0.29</b>	<b>-0.24</b>	<b>-0.28</b>
$\Delta y_t$	-0.20	-0.16	-0.19	-0.14	0.06	0.16
$\mathbb{E}_{\text{GB}}[\Delta y_{t+1}]$	0.02	0.13	0.16	<b>0.30</b>	0.07	0.08
Unemployment ( $u$ )	-0.23	-0.24	-0.10	-0.13	-0.20	<b>-0.35</b>
$\mathbb{E}_{\text{GB}}[u_{t+1}]$	-0.26	-0.28	-0.18	-0.24	-0.22	<b>-0.39</b>
$\pi_t^{\text{PCE}}$	<b>-0.61</b>	<b>-0.75</b>	<b>-0.68</b>	<b>-0.80</b>	<b>-0.41</b>	<b>-0.60</b>
$\mathbb{E}_{\text{GB}}[\pi_{t+1}]$	<b>-0.56</b>	<b>-0.70</b>	<b>-0.67</b>	<b>-0.78</b>	<b>-0.46</b>	<b>-0.57</b>

*Dependent Variable:* This table reports coefficients of univariate regressions of the three communications measures (transparency, and the entropies of the minutes and transcripts) on the variables in the rows of the table. The regressions are shown for the moving average of each measure, and the level of the measure. In all rows the relevant communication measures has been linearly detrended and standardized.

*Independent Variables:* The variable  $t$  is a linear trend. *TT State* is the “state of the economy” variable of Tenreyro and Thwaites (2016); “RR” are the monetary policy shocks of Romer and Romer (2004), updated through the sample; NS Policy Shock and NS Fed Funds Shock are the monetary and Fed-Funds futures shocks of Nakamura and Steinsson (2018); FFR is the federal funds target (or the midpoint of its target range, or the actual value when neither is available);  $\Delta y_t$  is the annualized quarterly growth rate of real GDP from FRED;  $u_t$  is the civilian unemployment rate from FRED;  $\pi_t^{\text{PCE}}$  is the annualized quarterly growth rate of the PCE price level; the rows  $\mathbb{E}_{\text{GB}}[x_{t+1}]$  correspond to the Greenbook forecast of  $x$  in the quarter following the FOMC meeting at time  $t$  (these correspond to the macroeconomic series from above, except that CPI inflation is used instead of PCE in order to have a longer sample from the Greenbook). The monetary policy shocks are timed such that the regression corresponds to the transparency of the minutes prevailing immediately before the shock is emitted. All variables have been standardized over the regression sample so that the coefficients reflect correlation coefficients.

*Sample:* The sample sizes [for the moving averages] are as follows: 74 for the NS policy shocks (Jan. 2004–Mar. 2014) to match the results from section 3.4; 276 [265] for the RR shocks (Mar. 1977–Oct. 2008); 261 for  $\mathbb{E}_{\text{GB}}[\pi_{t+1}]$  (Oct. 1979–Dec. 2014); and 326 [315] for all other variables (Mar. 1977–Dec.2014).

**Bolded estimates** are statistically significant at at least the 5% level, calculated using heteroskedasticity and autocorrelation-consistent asymptotic standard errors with the automatic lag selection method of Newey and West (1994), as implemented by Baum et al. (2010).

Table 3.3: Jointly Testing Correlates of Communications Measures

	Transparency		Minutes Entropy		Transcripts Entropy	
$t$	0.01**	(0.01)	0.01***	(0.00)	0.00	(0.01)
Burns (70–78)	2.24**	(1.06)	0.08	(0.41)	−0.78	(1.00)
Miller (78–79)	2.19**	(0.86)	−0.38	(0.32)	0.48	(1.01)
Volcker (79–87)	1.52**	(0.61)	−0.52**	(0.25)	−0.38	(0.70)
Greenspan (87–06)	1.05**	(0.42)	−0.42**	(0.20)	0.57	(0.43)
Bernanke (06–14)	0.05	(0.23)	−0.37***	(0.12)	0.48*	(0.27)
Yellen (14–18)	.	.	.	.	.	.
Broida (73–78)	.	.	.	.	.	.
Altmann (78–83)	0.07	(0.21)	−1.01***	(0.20)	0.76**	(0.37)
Axilrod (83–86)	0.03	(0.23)	0.61**	(0.25)	1.14***	(0.44)
Bernard (86–87)	0.22*	(0.12)	−0.06	(0.09)	0.30	(0.20)
Kohn (87–02)	0.13**	(0.06)	−0.12**	(0.06)	−0.11	(0.09)
Reinhart (02–07)	0.15	(0.19)	−0.30*	(0.16)	−0.97***	(0.35)
Madigan (07–10)	−0.48***	(0.16)	−0.58***	(0.13)	0.13	(0.30)
English (10–15)	−0.10	(0.19)	0.34**	(0.13)	−0.07	(0.28)
FFR	0.02	(0.04)	−0.03	(0.02)	−0.02	(0.03)
$\Delta$ FFR	−0.03	(0.04)	0.01	(0.02)	−0.04	(0.03)
$ \Delta$ FFR	−0.05	(0.04)	0.00	(0.00)	−0.01	(0.01)
Post-1993	0.33	(0.30)	−0.07	(0.23)	−0.58**	(0.27)
TT State	−0.39	(0.30)	−0.06	(0.20)	1.04***	(0.34)
$\Delta y_t$	0.13	(0.09)	−0.03	(0.02)	−0.03	(0.03)
Unemployment	−0.07	(0.07)	0.05*	(0.03)	0.09	(0.07)
$\pi_t^{\text{PCE}}$	0.01	(0.08)	0.02	(0.04)	0.08	(0.06)
Cons.	−3.00***	(1.13)	−0.46	(0.75)	−3.31*	(1.87)
<i>Chair</i>	0.000		0.000		0.000	
<i>Secretary</i>	0.001		0.000		0.000	
<i>Fed Funds</i>	0.632		0.050		0.577	
<i>Macro Vars.</i>	0.447		0.531		0.274	
$R^2$	0.935		0.868		0.954	
$N$	315		315		315	

Note: The table shows results for regressions of the three communications variables (in the columns) on dummies of the sitting chairs of the FOMC (Burns–Yellen); dummies for the secretaries (Broida–English); and several other variables defined in table 3.2. Chairs and secretaries have their years in office in parentheses. The rows with italicized labels contain p-values for tests of joint significance of groups of variables in the regression: *Chair* jointly tests the chair dummies; *Secretary* tests the secretary dummies; *Fed Funds* tests the three FFR variables; and *Macro Vars.* tests output growth, unemployment, and inflation. These p-values, as well as the standard errors in parentheses, are computed using asymptotic standard errors with the automatic lag selection method of Newey and West (1994), as implemented by Baum et al. (2010). The communications variables are standardized to be mean-zero with unit standard deviation over the regression sample, which is each regularly-scheduled FOMC meeting between March 15, 1977 and December 13, 2014 (the meetings between April 1976 and March 1977 are dropped when the moving average is formed).

table jump—albeit statistically insignificant—is that procedural transparency increased after the FOMC became aware that its transcripts would be released to the public in 1993, and the transcripts also became more focused (lower entropy). The fact that this event may have had an effect on the content of the transcripts—through its effect on the Committee’s deliberation—is the subject of a literature that has largely concluded (with the exception of the final paper) that the event caused a move towards less debate and more formal discussions (like what might be found in the minutes); see Meade (2005), Meade and Stasavage (2008), Acosta (2015), Egesdal et al. (2015), Hansen et al. (2017), and Woolley and Gardner (2017).

Next, the sitting Secretary of the FOMC is strongly correlated with the level of all three measures—the p-value for a test that these variables are jointly non-zero is negligible in all cases. This is encouraging, for it suggests that the person in charge of overseeing the creation of each document has a role in determining its properties. The sitting Chair is also correlated with all three measures. The Madigan dummy—a proxy for the Great Recession—in the minutes regression shows that the minutes became much more focused during the crisis and its aftermath.

By-and-large the measures are not predicted by current macroeconomic events and monetary policy once other controls are included, as can be seen in table 3.3. Unconditionally the variables tend to be countercyclical: transparency is lower when the Federal Funds rate and inflation are high, and both the minutes and transcripts becoming more focused during these times. The similarity between the minutes and the transcripts also tends to decrease in meetings preceding large (in absolute value) monetary policy shocks. In the case of the shocks of Nakamura and Steinsson (2018), this implies that markets are more surprised at time  $t + 1$  when the minutes from time  $t$  were less informative about the transcripts.

### 3.3.2 Anecdotal Understanding of the Measure


A less-systematic approach to understanding what the transparency measure is capturing is to read the underlying documents. Natural places to start are the points where the measure was particularly low and high. The lowest post-Volcker observation occurred in the December 1989

meeting, though the story starts in October 1989 with a meeting of the G7. In that meeting, the parties involved decided that the dollar was overvalued, and agreed to a coordinated action to flood the market with dollars in order to remedy this. In the October 1989 FOMC meeting, shortly after this coordinated action began, there was a debate about whether this action interfered with the Fed’s statutory mandate to achieve price stability. Some parties were concerned that the Fed would be “implicated in talking out of one side of [its] mouth about price stability goals and yet agreeing to constantly flooding the market with dollars” (Transcripts, 1976–2008, Governor Johnson). Others, like President Guffey, felt differently, stating “I’m not terribly concerned about the price stability issue in the sense that with sterilized intervention I think for some long period in the future we can go about a price stability objective without much problem” (Transcripts, 1976–2008, President Guffey). At the end of the discussion, Vice Chair Corrigan suggested that the staff prepare a “presentation for the Committee where [it] would take a look at this question of price stability in five years in some systematic way.” The December 1989 meeting contained that presentation and ensuing discussion, though there was no mention of it in the minutes of that meeting, this perhaps owing to President Guffey’s October concern about “bringing this issue to a confrontational stage outside the confines of this Committee and the Treasury.” This episode contrasts with the meeting of December 2004—the meeting with the highest level of transparency—in which the Committee undertook a lengthy discussion concerning communications policy. Specifically, they discussed the possibility of accelerating the release of the minutes from six to three weeks—a policy they subsequently implemented with the minutes of that meeting—and also made a record of this discussion in the minutes. These episodes highlight the fact that the measure captures meaningful discrepancies between the minutes and the transcripts.

### 3.3.3 Observability of the Measure

Lastly, figure 3.3 shows the transparency measure alongside a measure of transparency derived from newspapers. Specifically, in the spirit of Baker et al. (2016), Husted et al. (2020) construct a measure of monetary policy uncertainty by counting the number of articles in a given period of

Figure 3.3: Monetary Policy Transparency in the News



time that appear in major newspapers containing the terms “uncertainty,” “monetary policy,” and “Federal Reserve.”<sup>21</sup> The measure is divided by the number of articles that contain “Federal Reserve” for each newspaper in each period, in order to control for the volume of articles over time and the different focuses of each newspaper. After scaling the normalized counts by newspaper to have unit variance, the resulting series are summed to form the monetary policy uncertainty index. The blue line in figure 3.3 is constructed in nearly the same way, with the exclusion of “monetary policy” and “uncertainty” and the inclusion of “transparent” or “transparency.”<sup>22</sup> The resulting series is positively and significantly correlated with the transparency measure derived from the minutes and transcripts. Additionally, the moving average of the minute-transcript similarity is more-highly correlated with the newspaper-based measure (and its moving average) than the raw measure. This suggests that minute-transcript similarity, and its moving average, despite being based on a document that is not visible to the public, is something that is in some way observable, given its positive relationship with this clearly-observable newspaper-based measure. Perhaps this arises because much of the variation in transparency is driven by the sitting Chair, an

<sup>21</sup>They also include close synonyms of these three terms.

<sup>22</sup>A huge thanks to Lucas Husted, who constructed the new index.

easily observable variable.

### **3.4 Effectiveness of Monetary Policy**

Views on the role of transparency in monetary policy making have evolved greatly over the last 50 years. Today, transparency is often touted as a means through which monetary policy is made more effective—in 2013, then-Chair Ben Bernanke stated in a speech that “transparency about the framework of policy has aided the public in forming policy expectations, reduced uncertainty, and made policy more effective.” In this section I address this question—whether transparency makes monetary policy more effective—and provide evidence suggesting that it does.

#### **3.4.1 Defining Effectiveness**

The first question that arises when seeking an empirical answer to this question is how to define monetary policy effectiveness. One answer is a policy that allows a central bank to achieve its objective, such as price stability. Blinder et al. (2008) discuss the literature that has taken this approach. A general problem that arises is that establishing causal inference is challenging. Another possibility is that effective monetary policy is able to affect market expectations—Blinder et al. discuss the empirical literature that largely supports this proposition. In this vein, using a high-frequency identification strategy, Nakamura and Steinsson (2018) establish a causal link between monetary policy and real interest rates: the more-surprising the monetary policy announcement, the greater the movement in real interest rates.

While these empirical studies discussed in Blinder et al., Nakamura and Steinsson, and their predecessors have brought a deepened understanding of the effects of monetary policy and its communication, they say little about the role that transparency plays in determining these effects. As it pertains to transparency about the decisionmaking process, the answer is not obvious. An important characteristic of each meeting for understanding its transparency is the context in which it occurred—the state of the economy and the committee, for example. In a more-complicated policy

environment, when the content of the minutes are more-heavily scrutinized,<sup>23</sup> high transparency might increase or decrease uncertainty and, thus, the ability of policy to have any effects. Using cross-sectional variation in transparency Naszodi et al. (2016) highlight that transparency reduces forecast uncertainty, while increased volume of communication can have the opposite effect, as Lustenbergera and Rossib (2020) find. That said, previous work has postulated that increased transparency should enhance policymakers’ abilities to affect the real economy through longer term interest rates (e.g. Woodford, 2005), and the following exercises provide evidence supporting this hypothesis.

### 3.4.2 Empirical Strategy

My empirical strategy builds off of the work of Nakamura and Steinsson, who estimate the following equation:

$$\Delta i_t = a + b \varepsilon_t + \text{error}_t \quad (3.1)$$

where  $i_t$  stands for a multitude of nominal and real interest rate forwards and yields. I take estimates of daily nominal rates from Gürkaynak et al. (2007), and real rates from Gürkaynak et al. (2010).<sup>24</sup> I start my analysis in 2004 since that is when data on all interest rates under consideration are available, as discussed in Gürkaynak et al. (2010). The variable  $\varepsilon_t$  is a high-frequency monetary policy shock identified as the first principal component of the change in Fed-Funds and Eurodollar futures out to four quarters.<sup>25</sup> This change is taken over a narrow window around FOMC statement releases. Nakamura and Steinsson find that, in response to a monetary shock, nominal and real interest rates move by a similar amount several years into the term structure. The fact that *real* interest rates move is taken as evidence of monetary non-neutrality.

---

<sup>23</sup>A simple Google Trends search confirms this—“Fed Meeting” and “Fed Minutes” were the most popular around late 2007, mid 2013, and late 2015; corresponding the beginning of a monetary easing, the months the “Taper Tantrum,” and the departure from the zero lower bound.

<sup>24</sup>These are available for download from the FEDS working-paper version of these papers: Gürkaynak et al. (2006) and Gürkaynak et al. (2008) respectively.

<sup>25</sup>I take this from the replication materials of Nakamura and Steinsson.

### 3.4.3 Results and Robustness

I slightly augment equation (3.1) in order to answer the question of whether monetary policy is more or less effective when the Fed is more-accurately representing the content of its meetings:

$$\Delta i_t = \alpha + \beta \varepsilon_t + \gamma (\varepsilon_t \times \bar{\tau}_{t-1}) + \phi \bar{\tau}_{t-1} + \text{error}_t, \quad (3.2)$$

where  $\bar{\tau}_t$  is the twelve-meeting moving average of minute-transcript similarity, which has been standardized for interpretability.<sup>26</sup> Because the minutes of meeting  $t - 1$  are released between  $t - 1$  and  $t$ , the value of  $\bar{\tau}$  at time  $t - 1$  reflects the prevailing level of procedural transparency at time  $t$ . The coefficient of interest is  $\gamma$ . With  $\gamma > 0$ , monetary policy has larger effects on interest rates when transparency is above its mean—a standardization of  $\bar{\tau}_t$  means that  $\gamma$  can be interpreted as the interaction effect when procedural transparency is one standard deviation above its mean.

Figure 3.4 shows the estimated coefficients.<sup>27</sup> The estimated values of  $\gamma$  are positive for nearly every interest rate under consideration, with several of these being statistically significantly different than zero. Averaging the estimates of  $\beta + \gamma$  for real yields shows that the effect of monetary policy shocks on real yields is 43 percent higher when transparency is one standard deviation above its mean than when it is at its average level. The estimates of  $\beta$  and  $b$  are consistent with the estimates and conclusion of Nakamura and Steinsson (2018)—that nominal and real rates move together far out into the term structure—even using a slightly different sample period.

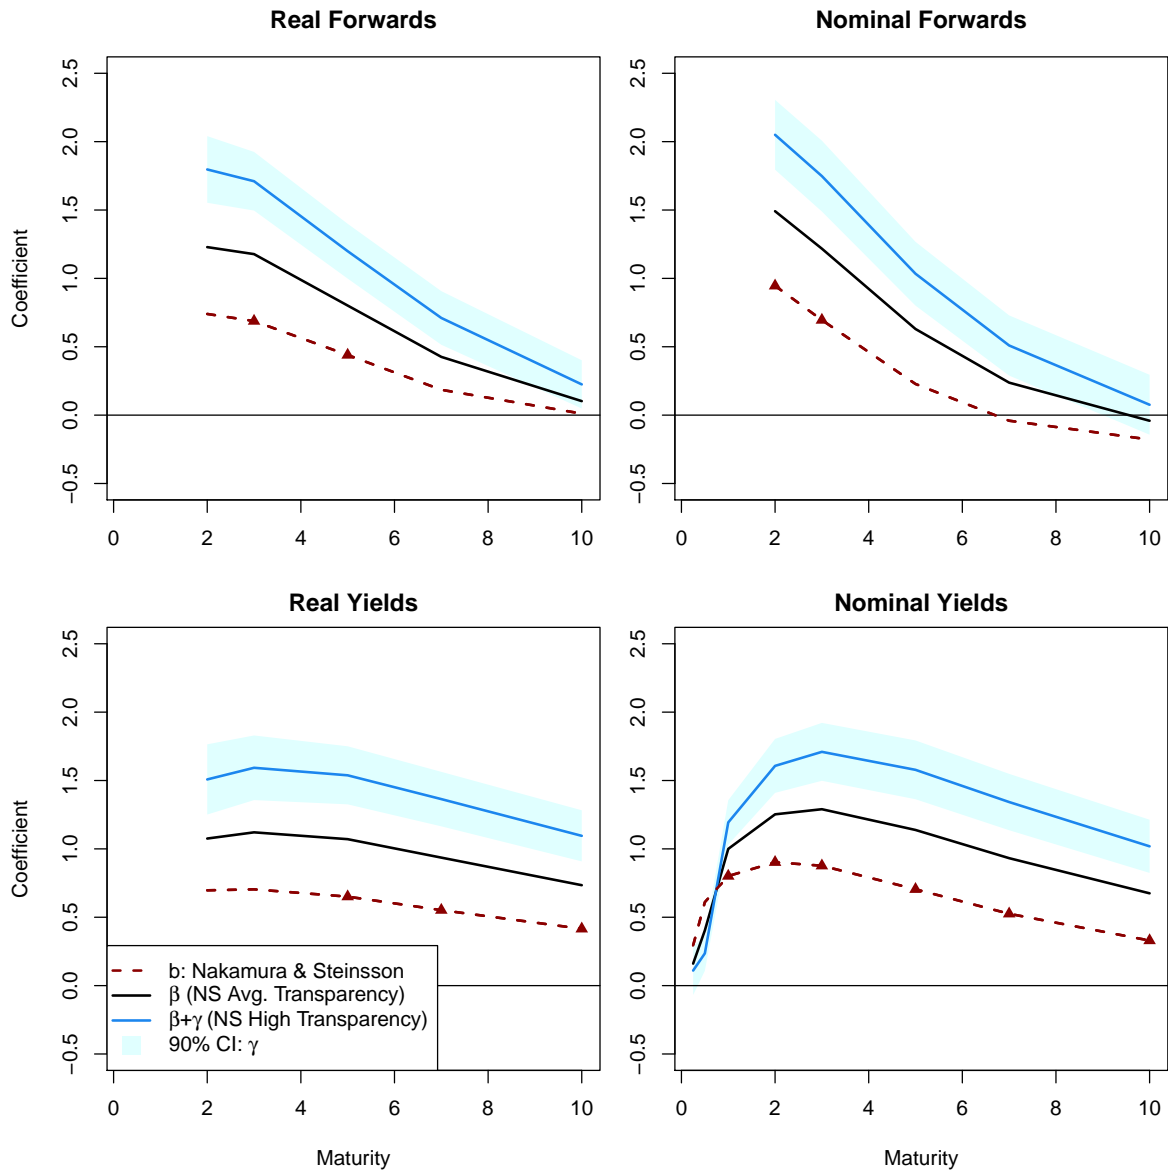
The discussion in section 3.3 leaned against an interpretation of minute-transcript similarity as a proxy for something unrelated to procedural transparency. That said, one might be worried that the correlation of transparency with the entropy of the transcripts (which was shown to be

---

<sup>26</sup>The reason for using a moving average has been discussed above. Twelve meetings is one-year worth of meetings at the beginning of the sample—this switches to eight meetings in 1981. The results are robust to using several other lag lengths—see Q. Section 3.3 pointed to noise and observability as reasons motivating the use of a moving average—the latter is important here. If this procedural transparency is not perceptible by the public then it is difficult to imagine reasons for which it might have effects. Given that transcripts are not released for several years, it seems unlikely that the public could realize the level of procedural transparency in real time. However, given that the Committee members give speeches and other public commentary, the public should have a sense of what is on the mind of Committee members and be able to compare this to the concerns enumerated in the minutes.

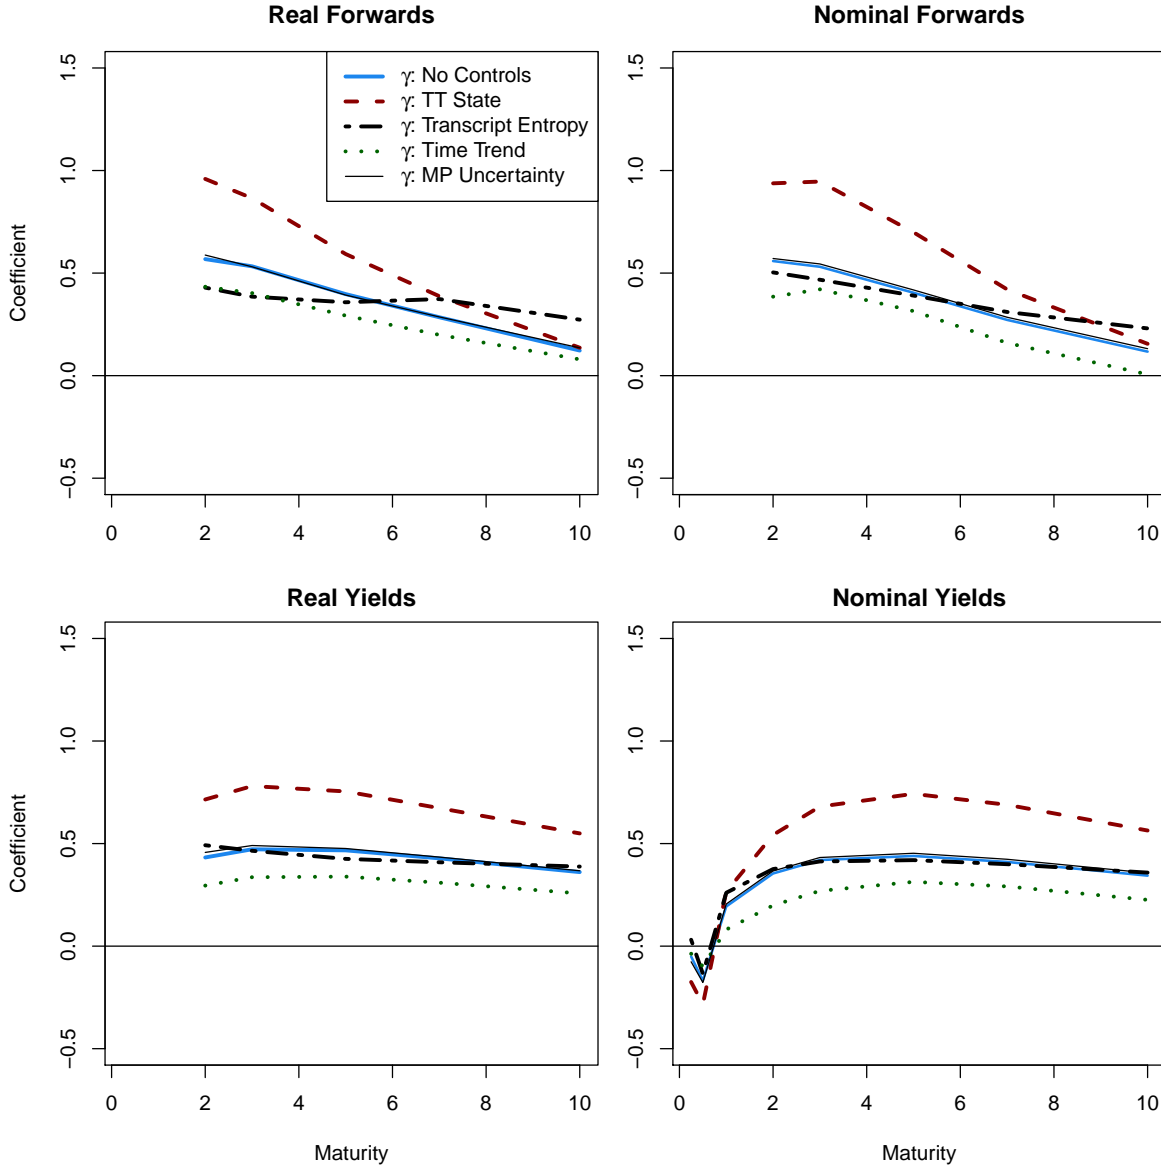
<sup>27</sup>Appendix R contains a table of the estimates.

Figure 3.4: Monetary Policy Effectiveness: Regression Results



Note: These figures show the results of estimating equations 3.1 and 3.2. The four panels correspond to whether  $i_t$  is a real or nominal interest rate yield or forward. The x-axis in each plot refers to the relevant maturity for each rate (3- and 6-month, and 1-, 2-, 3-, 5-, 7- and 10-year rates). The blue lines show  $\beta + \gamma$  surrounded by the 90% confidence interval of  $\gamma$  computed using heteroskedasticity-robust standard errors. The black lines show  $\beta$ , and the red-dashed lines show  $b$ — $b$  will be different than the corresponding values in Nakamura and Steinsson (2018), since my sample runs from January 2004–March 2014 (in contrast to their sample of 2000– March 2014) though I follow Nakamura and Steinsson in dropping the July 2008 – July 2009 period (which has little impact on the results). This forms 74 observations. The shocks are scaled so that the effect of the shock on one-year nominal yields is unity when transparency is at its average level. Transparency is standardized to have unit variance and zero-mean over the sample, so that  $\gamma$  corresponds to the interaction effect when transparency is one standard deviation above its mean. The red triangles are drawn whenever the estimated value of  $b$  is statistically significantly below  $\beta$  at at-least the 10% level, again using robust standard errors.

Figure 3.5: Monetary Policy Effectiveness: Robustness of Regression Results



Note: The graphs show the coefficients  $\gamma$  estimated based on equation (3.3). See the note to figure 3.4 for details about the sample. The lines labeled “TT State” include the “state of the economy” variable of Tenreyro and Thwaites (2016) as  $x_t$ , where  $t$  refers to the quarter in which the FOMC meeting took place. The lines labeled “Transcript Entropy” have  $x_t = H(\theta_{t-1})$ —the lagged entropy of the transcripts. The “Time Trend” label refers to the case in which  $x_t = t$ . Finally, “MP Uncertainty” is the level of uncertainty about monetary policy—as measured by Husted et al. (2020)—as it stood at the end of the previous FOMC meeting. Again, the monetary policy shocks are normalized so that  $\beta = 1$  for the one-year nominal yields regression over the sample period.

correlated with economic conditions) would lead the estimates of equation (3.2) to simply replicate the conclusions of Tenreyro and Thwaites (2016)—namely, that monetary policy is less effective during recessions. One might also worry that transparency—having a slightly positive trend over time—is serving as a proxy for a Fed that has increasingly relied on longer-term forward guidance. Another concern is that this measure could be proxying for uncertainty regarding monetary policy. In order to alleviate these worries, figure 3.5 shows the estimates of  $\gamma$  when different controls—and their interactions with the monetary shock—have been included in the estimating equation:

$$\begin{aligned} \Delta i_t = & \alpha + \beta \varepsilon_t + \gamma(\varepsilon_t \times \bar{\tau}_{t-1}) + \phi \bar{\tau}_{t-1} \\ & + \omega(\varepsilon_t \times x_t) + \pi x_t + \text{error}_t, \end{aligned} \quad (3.3)$$

where  $x_t$  is the entropy of the transcripts, the state variable of Tenreyro and Thwaites (for the first concern), a time trend (for the second concern), or the monetary policy uncertainty index of Husted et al. (2020). The results are consistent with the earlier findings— $\gamma$  is positive for nearly every interest rate under consideration.

#### 3.4.4 Omitted Variables Bias

Finally, figure 3.4 also highlights a slight downward bias when equation (3.1) is estimated without controlling for the role that transparency plays. It is typically the case that  $b < \beta$ , and in a few cases this difference is statistically significant—red triangles are shown whenever  $b < \beta$  is statistically significant at at least the 10% level.<sup>28</sup> Why is this the case? Mechanically, as in any omitted variables bias problem, one has to consider the following relationship:

$$\varepsilon_t \times \bar{\tau}_{t-1} = \psi_0 + \psi \varepsilon_t + \text{error}_t.$$

The estimate of  $\psi$  will have the sign of  $\text{cov}(\varepsilon_t \times \bar{\tau}_{t-1}, \varepsilon_t) \approx \mathbb{E}[\varepsilon_t^2 \bar{\tau}_{t-1}]$ , which was shown to be slightly negative in table 3.2—that is, monetary policy shocks tend to be larger when transparency

---

<sup>28</sup>This is computed using a seemingly-unrelated regressions model.

is lower. With this relationship, however, the estimate of  $b$  will not be  $b$  but instead  $\hat{b} = \beta + \gamma \cdot \psi$ . With  $\psi < 0$  and  $\gamma > 0$ , this implies that  $b$  underestimates the true effect of monetary policy shocks,  $\beta$ . The intuition for this result is as follows. Consider a large positive shock. The largest monetary shocks tend to occur when transparency is low ( $\psi < 0$ ). However, low transparency also means that the effect of these large shocks on interest rates will be lower (the interaction effect,  $\gamma \varepsilon_t \bar{\tau}_{t-1}$ , is small or negative on average with small or negative  $\bar{\tau}_{t-1}$ ). Theoretical work may help to clarify this chain of events, though it does suggest that transparency may be a double-edged sword for the effectiveness of monetary policy, if moving interest rates is the definition of effectiveness. Larger monetary shocks are emitted when the Fed is being less transparent about its discussions. But, possibly because the public cannot make as much sense of these short-run surprises, the shocks are less-easily transmitted to longer term interest rates. The next section provides empirical evidence in support of this interpretation.

### **3.5 Interpretation: Transparency about What?**

The main transparency measure used in the estimation of equation (3.2) is an aggregate measure of transparency, though the richness of the underlying data allows for transparency to be measured along different dimensions. In section 3.5.1, I revisit my estimates using two new measures of transparency: transparency about discussion of the economic outlook, and transparency about monetary policy strategies. I find that my empirical estimates are driven by the latter. To provide additional color to these results I estimate in section 3.5.2 transparency regarding topics that can be used to “predict” inflation and FOMC dissents, described at end of section 3.2.3. Consistent with the results in section 3.5.1, I find that transparency about monetary policy discussions—presumably contentions discussions if they predict FOMC dissents—allow interest rate shocks to pass through more fully to longer term rates. That is not robustly the case for transparency about topics that primarily reveal information about the state of the economy.

### 3.5.1 Economic Outlook and Policy Strategy Transparency

The structure of the transcripts suggests a natural first-step towards understanding what minute-transcript similarity is picking up, and which aspect of it plays an important role in determining the effectiveness of monetary policy. As documented by Hansen et al. (2017), “FOMC meetings have two major parts related to the monetary policy decision: the economic situation discussion... followed by the monetary policy strategy discussion.” The authors treat these sections as separate in their analysis, and in this section I follow their lead—in so doing, I refer to the first section as ECSIT, and the latter as MPS. I use the breakdowns of the transcript and minutes in order to create an ECSIT transparency index, and an MPS transparency index.<sup>29</sup> The ECSIT index reflects the extent to which the Committee provides details about its reading of the state of the economy to the public through the minutes. The MPS index, on the other hand, captures discussions about what this reading implies for monetary policy. The MPS therefore include discussions of the models and targets preferred by policymakers, and their policy preferences more generally.

In order to create these indexes, I estimate the document distributions for the ECSIT and MPS sections of the minutes and transcripts of each meeting.<sup>30</sup> This gives, for every FOMC meeting, a measurement of the transparency of the ECSIT portion of the meeting, and of the MPS portion, displayed in figure 3.6.


Figures 3.7 and 3.8 repeat the analysis regarding the role that transparency plays in determining the effectiveness of monetary policy. Specifically, they present the results of estimating equations 3.1 and 3.2, where  $\bar{\tau}_{t-1}$  is replaced with the twelve-meeting moving-average of MPS and ECSIT similarity, respectively. The results for MPS similarity are quite similar to those that included minute-transcript similarity (figure 3.7). This is not true for ECSIT similarity, despite the fact that ECSIT takes up a much-larger portion of the FOMC discussion (roughly 40% vs. 20 % for MPS).

---

<sup>29</sup>Stephen Hansen graciously supplied the breakdown of the transcripts through 2011, and I updated this breakdown through 2014. I also performed the corresponding split for the minutes manually since 1995. To have a sense, the ECSIT portion of the minutes in 2014 included the sections titled *Staff Review of the Economic Situation*, *Staff Review of the Financial Situation*, *Staff Economic Outlook*, and *Participants’ Views on Current Conditions and the Economic Outlook*, and the MPS section was the section titled *Committee Policy Action*. More details can be provided on request.

<sup>30</sup>Using the same estimated LDA topics as above.

Figure 3.6: Transparency by Part of Transcript



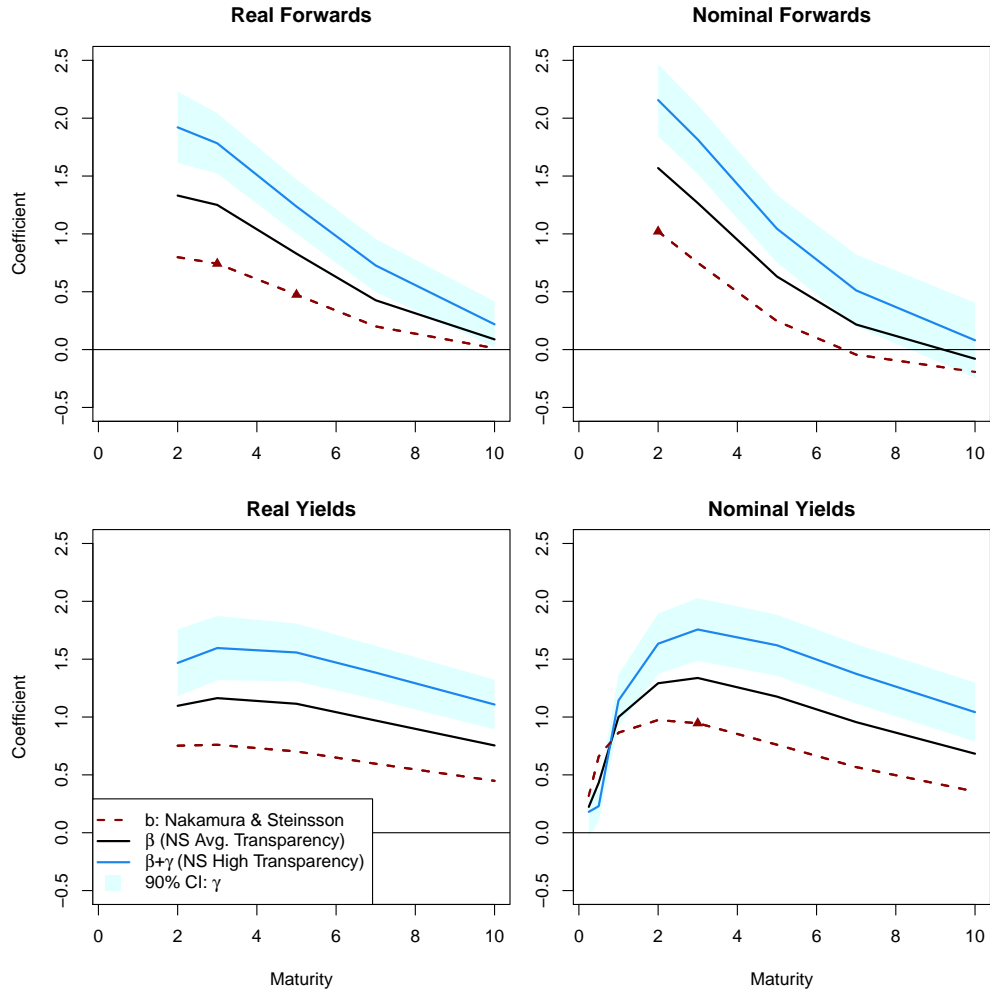
Note: This graph shows the 12-meeting moving-average of transparency using the full minutes and transcripts (solid black line); the ECSIT portions of the minutes and transcripts (red-dashed line); and the MPS portion of the minutes and transcripts (blue-dotted lines). The bracketed numbers represent the correlation coefficient of the (moving average of the) ECSIT and MPS transparency measures with the full-document transparency measure. For reference, the ECSIT and MPS discussions take up on average 73 and 12 percent of the minutes, and 51 and 25 percent of the transcripts. These ratios have remained essentially constant since 1995.

Thus, the fact that monetary policy is more effective when transparency is elevated owes more to transparency regarding policymakers' views about the appropriate monetary policy—conditional on their reading of the state of the economy—than to transparency about the readings themselves.

### 3.5.2 Topic-Specific Transparency

An alternative approach for measuring transparency along different dimensions is to focus on the transparency of specific topics estimated by the LDA language model. In this section I study the role of transparency about topics that are similar to those presented in the previous section: transparency about monetary policy, and transparency about economic fundamentals. To assess the first, I measure the transparency about topics whose presence in the minutes or transcripts predict the number of dissents at each FOMC meeting in a LASSO regression. Of the 50 topics estimated by LDA, I denote the topics selected to predict dissents by  $\mathcal{K}^D \subseteq \{1, \dots, 50\}$ . For transparency regarding economic fundamentals, I measure the transparency about topics whose

Figure 3.7: Monetary Policy Effectiveness: Regression Results using MPS Similarity

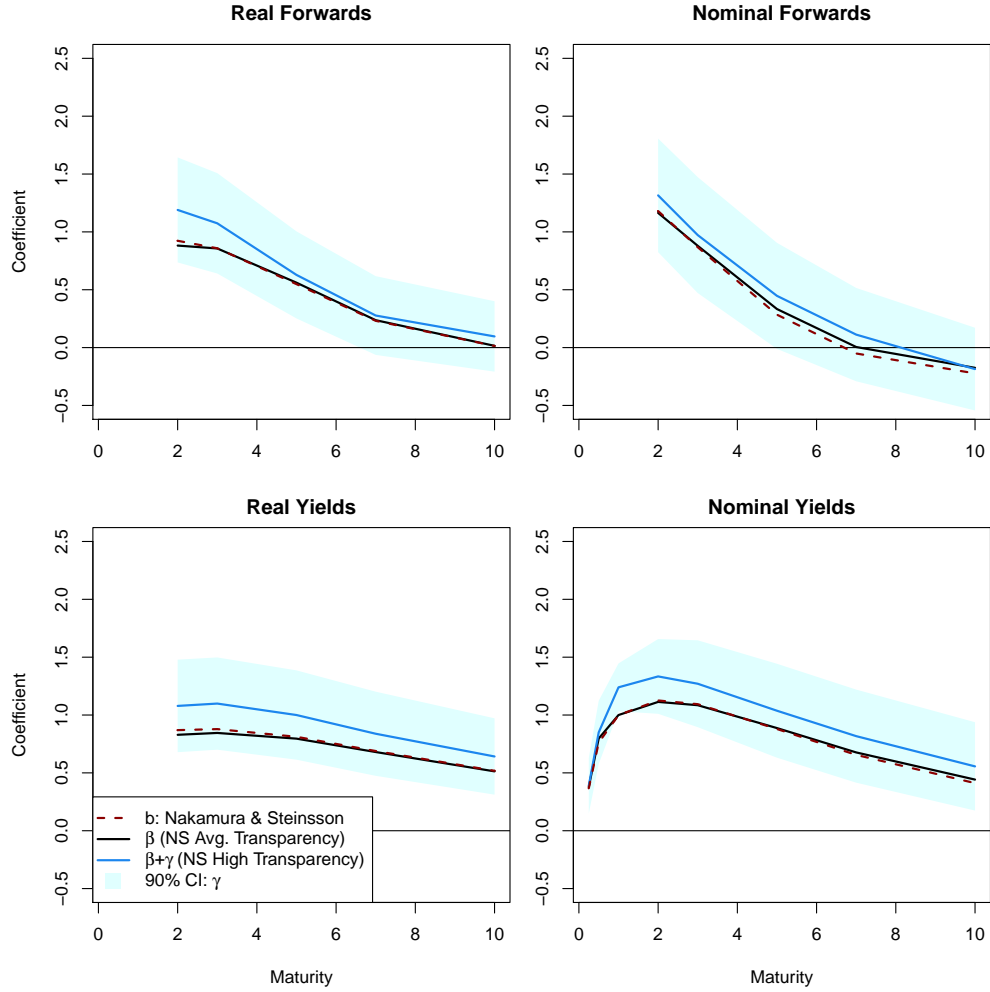


Note: These graphs present the results of estimating equations 3.1 and 3.2, where  $\bar{\tau}_{t-1}$  is replaced with the twelve-meeting moving-average of MPS similarity. Everything else is the same as in figure 3.4, so its note can be referenced for further details.

presence in the minutes or transcripts predict the level of CPI inflation in the month corresponding to each FOMC meeting. I denote these topics by  $\mathcal{K}^\pi$ . More details regarding the selection of these topics can be found at the end of section 3.2.3.

I define the transparency for dissent-related topics  $\tau_t^D$  and inflation-related topics  $\tau_t^\pi$  as the cosine similarity between the (truncated) distributions of the minutes and transcripts over the selected topics. Recalling from section 3.2.4 that the topic distributions of the minutes and transcripts are

Figure 3.8: Monetary Policy Effectiveness: Regression Results using ECSIT Similarity



Note: These graphs present the results of estimating equations 3.1 and 3.2, where  $\bar{\tau}_{t-1}$  is replaced with the twelve-meeting moving-average of ECSIT similarity. Everything else is the same as in figure 3.4, so its note can be referenced for further details.

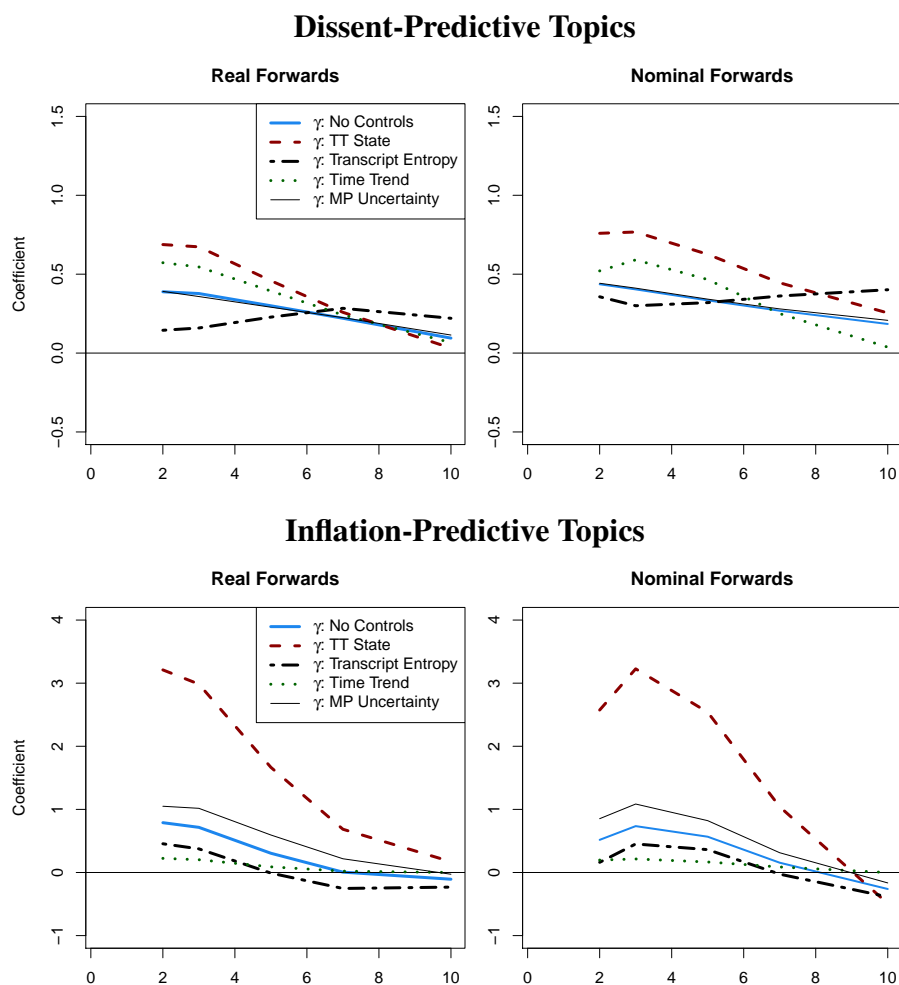
given by  $\mu_t$  and  $\theta_t$ , the transparency measures are given by

$$\tau_t^D = \frac{\sum_{k \in \mathcal{K}^D} \theta_t^k \mu_t^k}{\sqrt{(\sum_{k \in \mathcal{K}^D} (\theta_t^k)^2) (\sum_{k \in \mathcal{K}^D} (\mu_t^k)^2)}}$$

$$\tau_t^\pi = \frac{\sum_{k \in \mathcal{K}^\pi} \theta_t^k \mu_t^k}{\sqrt{(\sum_{k \in \mathcal{K}^\pi} (\theta_t^k)^2) (\sum_{k \in \mathcal{K}^\pi} (\mu_t^k)^2)}}$$

Cosine similarities are a commonly used measure of document similarity in the NLP literature, and

Figure 3.9: Monetary Policy Effectiveness: Regression Results using Topics-Specific Similarities



Note: The graphs show the coefficients  $\gamma$  estimated based on equation (3.3), where  $\bar{\tau}_{t-1}$  is replaced with the twelve-meeting moving-average of the similarities of the topics in the minutes and transcripts that are predictive of the number of FOMC dissents and inflation. Everything else is the same as in figure 3.5, so its note can be referenced for further details. Notice the wider axis for the Inflation panel.

are simply the uncentered correlation coefficient between the topics used in each document.<sup>31</sup>

Figure 3.9 shows the estimates of  $\gamma$  for the two transparency measures for specifications that include the various controls considered in the construction of figure 3.5. The estimates of  $\gamma$  from the specification that uses  $\tau^\pi$  are not robustly positive, in contrast to estimates of  $\gamma$  estimated using dissent-predictive topics. This further suggests that it is topics relevant the setting of monetary

<sup>31</sup>I use cosine similarity instead of the Kullback-Leibler divergence because the distributions of the minutes and transcripts over the selected topics do not sum to unity.

policy—conditional on an economic outlook—that causes changes in short-term interest rate to affect longer term nominal and real interest rates.

### **3.6 Conclusion**

A central bank has a plethora of channels through which it can be transparent. Whether it is the publication of inflation reports, timely summaries of policy decisions, or post-meeting press conferences, the objective typically is to explain to the public the rationale behind policy decisions. Nowhere can these rationales be better captured than in the actual meetings in which these decisions are considered and made. The measure of transparency I propose in this paper—the similarity between the minutes and transcripts of each FOMC meeting—captures the extent to which the content of these meetings is described to the public. While this minute-transcript similarity has fluctuated over time, generally the Fed has become more transparent about its reasoning over the last 40 years. My measure is only weakly correlated with economic conditions and the policy-making environment more generally, further supporting the case that minute-transcript similarity captures something more than simply fluctuations in the discussions of FOMC meetings. Anecdotal evidence shows that this measure does indeed capture meaningful discrepancies between the two documents, and there is little evidence to suggest that these discrepancies are purposeful—they are more likely to do with the fact that the writers of the minutes face a difficult task.

Evidence suggests that monetary policy shocks have larger effects on interest rates when minute-transcript similarity is high. Additionally, when the role of transparency is neglected, these shocks seem to have smaller effects because the largest of these shocks tend to be delivered at times when transparency is low, which is also when the shocks have smaller effects on interest rates. These results suggest that high transparency allows the public to better understand what monetary policy communications and short-term interest rate movements imply for the path of future policy, captured by longer-term interest rates.

## References

- Acosta, Miguel**, “FOMC Responses to Calls for Transparency,” *Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series*, 2015, 60.
- **and Hassan Afrouzi**, “Rationally Confused: On the Aggregate Implications of Information Provision Policies,” Technical Report In Progress.
- **and Joseph M. Saia**, “Estimating the Effects of Monetary Policy via High Frequency Factors,” Technical Report In progress.
- Andrade, Philippe and Filippo Ferroni**, “Delphic and Odyssean monetary policy shocks: Evidence from the euro area,” *Journal of Monetary Economics*, 2021, 117, 816–832.
- Angrist, Joshua D and Jörn-Steffen Pischke**, *Mostly harmless econometrics: An empiricist’s companion*, Princeton university press, 2008.
- Arai, Natsuki**, “The effects of monetary policy announcements at the zero lower bound,” *International Journal of Central Banking*, 2017, 13 (2), 159–196.
- Bagwell, Kyle, Robert W Staiger, and Ali Yurukoglu**, “Multilateral trade bargaining: A first look at the GATT bargaining records,” *American Economic Journal: Applied Economics*, 2020, 12 (3), 72–105.
- Baker, Scott R., Nicholas Bloom, and Steven J. Davis**, “Measuring Economic Policy Uncertainty,” *The Quarterly Journal of Economics*, 2016, 131 (4), 1593–1636.
- Bauer, Michael D. and Eric T. Swanson**, “The Fed’s Response to Economic News Explains the ‘Fed Information Effect’,” Working Paper 2020.
- Baum, Christopher F, Mark E Schaffer, and Steven Stillman**, “ivreg2: Stata module for extended instrumental variables/2SLS, GMM and AC/HAC, LIML and k-class regression.” 2010.
- Bernanke, Ben S.**, “Central Bank Independence, Transparency, and Accountability,” in “Speech at the Institute for Monetary and Economic Studies International Conference, Bank of Japan, Tokyo, Japan” 2010.
- , “Communication and monetary policy,” in “Speech at the National Economists Club Annual Dinner, Herbert Stein Memorial Lecture, Washington, DC” 2013.
- Binder, Carola**, “Fed speak on main street: Central bank communication and household expectations,” *Journal of Macroeconomics*, 2017, 52, 238–251.

- Blanchard, Olivier J., Jean-Paul L’Huillier, and Guido Lorenzoni**, “News, Noise, and Fluctuations: An Empirical Exploration,” *American Economic Review*, December 2013, 103 (7), 3045–70.
- Blei, David M, Andrew Y Ng, and Michael I Jordan**, “Latent Dirichlet Allocation,” *Journal of Machine Learning Research*, 2003, 3 (Jan), 993–1022.
- Blinder, Alan S., Michael Ehrmann, Marcel Fratzscher, Jakob De Haan, and David Jan Jansen**, “Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence,” *Journal of Economic Literature*, 2008, 46 (4), 910–945.
- Boguth, Oliver, Vincent Gregoire, and Charles Martineau**, “Information Flow, Noise, and the Irrelevance of FOMC Announcement Returns,” 2022.
- Borusyak, Kirill and Xavier Jaravel**, “The Distributional Effects of Trade: Theory and Evidence from the United States,” Working Paper 2021.
- Bown, Chad P.**, *Self-Enforcing Trade: Developing Countries and WTO Dispute Settlement*, Brookings Institution Press, 2009.
- Bown, Chad P and Douglas A Irwin**, “The GATT’s Starting Point: Tariff Levels circa 1947,” Technical Report, National Bureau of Economic Research 2015.
- and —, *The GATT’s Starting Point: Tariff Levels Circa 1947*, The World Bank, 2016.
- Brave, Scott A, R Andrew Butters, David Kelley et al.**, “A new “big data” index of US economic activity,” *Economic Perspectives, Federal Reserve Bank of Chicago*, 2019, 1.
- Brossard, E. B.**, “Comparison of Rates of Duty in the Pending Tariff Bill of 1929 with the Tariff Act of 1913 and the Tariff Act of 1922: Letter from USTC Chairman Brossard to Senator Reed Smoot,” Mar 1930.
- Cai, Yong, Santiago Camara, and Nicholas Capel**, “It’s not always about the money, sometimes it’s about sending a message: Evidence of Informational Content in Monetary Policy Announcements,” Technical Report, Working Paper 2021.
- Campbell, Jeffrey R, Charles L Evans, Jonas D.M. Fisher, and Alejandro Justiniano**, “Macroeconomic effects of Federal Reserve forward guidance,” *Brookings Papers on Economic Activity*, 2012, pp. 1–80.
- Cavallo, Alberto and Roberto Rigobon**, “The billion prices project: Using online prices for measurement and research,” *Journal of Economic Perspectives*, 2016, 30 (2), 151–78.
- Chahrour, Ryan and Kyle Jurado**, “Recoverability and expectations-driven fluctuations,” *Review of Economic Studies. Forthcoming*, 2021.

- Cieslak, Anna and Andreas Schrimpf**, “Non-monetary news in central bank communication,” *Journal of International Economics*, 2019, 118, 293 – 315.
- Clarkson, Douglas B**, “Estimating the standard errors of rotated factor loadings by jackknifing,” *Psychometrika*, 1979, 44 (3), 297–314.
- Coibion, Olivier**, “Are the effects of monetary policy shocks big or small?,” *American Economic Journal: Macroeconomics*, 2012, 4 (2), 1–32.
- , **Yuriy Gorodnichenko, and Michael Weber**, “Monetary Policy Communications and their Effects on Household Inflation Expectations,” *Journal of Political Economy*, Forthcoming.
- Crucini, Mario J**, “Sources of variation in real tariff rates: The United States, 1900-1940,” *The American Economic Review*, 1994, 84 (3), 732–743.
- D’Amico, Stefania and Thomas King**, “What Does Anticipated Monetary Policy Do?,” Technical Report, Working Paper 2017.
- Danker, Deborah J. and Matthew M. Luecke**, “Background on FOMC Meeting Minutes,” *Federal Reserve Bulletin*, 2005, 91, 175.
- Dincer, N. Nergiz and Barry Eichengreen**, “Central Bank Transparency and Independence: Updates and New Measures,” *International Journal of Central Banking*, 2014, 10 (1), 189–259.
- Doh, Taeyoung, Dongho Song, and Shu-Kuei Yang**, “Deciphering Federal Reserve communication via text analysis of alternative FOMC statements,” Technical Report, Working Paper 2020.
- Egedal, Michael, Michael Gill, and Martin Rotemberg**, “How Federal Reserve Discussions Respond to Increased Transparency,” *Available at SSRN*, 2015.
- Eijffinger, Sylvester C.W. and Petra M. Geraats**, “How Transparent are Central Banks?,” *European Journal of Political Economy*, 2006, 22 (1), 1–21.
- Ellen, Saskia Ter, Vegard H Larsen, and Leif Anders Thorsrud**, “Narrative monetary policy surprises and the media,” *Journal of Money, Credit, and Banking*, forthcoming.
- Fajgelbaum, Pablo D and Amit K Khandelwal**, “Measuring the unequal gains from trade,” *The Quarterly Journal of Economics*, 2016, 131 (3), 1113–1180.
- Feenstra, Robert C**, “US imports, 1972-1994: Data and concordances,” Technical Report, National Bureau of Economic Research 1996.
- , **John Romalis, and Peter K Schott**, “U.S. Imports, Exports, and Tariff Data, 1989-2001,” Working Paper 9387, National Bureau of Economic Research December 2002.

- Fernández-Villaverde, Jesús, Juan F. Rubio-Ramírez, Thomas J. Sargent, and Mark W. Watson**, “ABCs (and Ds) of Understanding VARs,” *American Economic Review*, June 2007, 97 (3), 1021–1026.
- Frisch, Ragnar and Frederick V Waugh**, “Partial time regressions as compared with individual trends,” *Econometrica: Journal of the Econometric Society*, 1933, pp. 387–401.
- Furman, Jason, Kathryn Russ, and Jay Shambaugh**, “US tariffs are an arbitrary and regressive tax,” *VoxEU.org* (Center for Economic and Policy Research blog). <http://voxeu.org/article/us-tariffs-are-arbitrary-and-regressive-tax> (January 15, 2017), 2017.
- Galí, Jordi**, *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*, Princeton University Press, 2015.
- Geraats, Petra M**, “Why adopt transparency? The publication of central bank forecasts,” 2001.
- , “Central Bank Transparency,” *The Economic Journal*, 2002, 112 (483), F532–F565.
- , “The Mystique of Central Bank Speak,” *International Journal of Central Banking*, 2007, 3 (1), 37–80.
- Gertler, Mark and Peter Karadi**, “Monetary Policy Surprises, Credit Costs, and Economic Activity,” *American Economic Journal: Macroeconomics*, January 2015, 7 (1), 44–76.
- Goodfriend, Marvin**, “Monetary Mystique: Secrecy and Central Banking,” *Journal of Monetary Economics*, 1986, 17 (1), 63–92.
- Gowa, Joanne and Raymond Hicks**, ““Big” Treaties, Small Effects: The RTAA Agreements,” *World Politics*, 2018, 70 (2), 165–193.
- Grant, Matthew**, “Where do the Data Come From? Endogenous Classification in Administrative Data,” 2021.
- Gresser, Edward**, “Toughest on the Poor: America’s Flawed Tariff System,” *Foreign Affairs*, 2003, 81 (6), 9–14.
- Griffiths, Thomas L and Mark Steyvers**, “Finding scientific topics,” *Proceedings of the National Academy of Sciences*, 2004, 101 (suppl 1), 5228–5235.
- Gulotty, Robert**, “Mismeasured Cooperation: The Trans-Atlantic Origins of Reciprocity in Market Access,” 2018.
- Gürkaynak, Refet S., Brian P. Sack, and Eric T. Swanson**, “Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements,” *International Journal of Central Banking*, 2005.

- , – , and – , “Market-based measures of monetary policy expectations,” *Journal of Business & Economic Statistics*, 2007, 25 (2), 201–212.
- , – , and **Jonathan H. Wright**, “The U.S. Treasury Yield Curve: 1961 to the Present,” *Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series*, 2006, 28.
- , – , and – , “The TIPS Yield Curve and Inflation Compensation,” *Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series*, 2008, 5.
- , – , and – , “The TIPS yield curve and inflation compensation,” *American Economic Journal: Macroeconomics*, 2010, 2 (1), 70–92.
- , **Burçin Kisacikoğlu**, and **Jonathan H. Wright**, “Missing Events in Event Studies: Identifying the Effects of Partially Measured News Surprises,” *American Economic Review*, 2020, 110 (12), 3871–3912.
- Handlan, Amy**, “Text Shocks and Monetary Surprises: Text Analysis of FOMC Statements with Machine Learning,” Working Paper 2020.
- Hansen, Lars Peter**, “Large sample properties of generalized method of moments estimators,” *Econometrica: Journal of the econometric society*, 1982, pp. 1029–1054.
- Hansen, Stephen and Michael McMahon**, “Shocking language: Understanding the Macroeconomic Effects of Central Bank Communication,” *Journal of International Economics*, 2016, 99.
- , – , and **Andrea Prat**, “Transparency and Deliberation within the FOMC: a Computational Linguistics Approach,” *Quarterly Journal of Economics*, 2017.
- , – , and **Matthew Tong**, “The long-run information effect of central bank communication,” *Journal of Monetary Economics*, 2019, 108, 185–202.
- Haubrich, Joseph, George Pennacchi, and Peter Ritchken**, “Inflation expectations, real rates, and risk premia: Evidence from inflation swaps,” *The Review of Financial Studies*, 2012, 25 (5), 1588–1629.
- He, Zheli**, “Trade and Real Wages of the Rich and Poor: Cross-Region Evidence,” 2018.
- Hébert, Benjamin and Jesse Schreger**, “The costs of sovereign default: Evidence from argentina,” *American Economic Review*, 2017, 107 (10), 3119–45.
- Hoesch, Lukas, Barbara Rossi, and Tatevik Sekhposyan**, “Has the information channel of monetary policy disappeared? Revisiting the empirical evidence,” Working Paper 2020.
- Hottman, Colin J and Ryan Monarch**, “A matter of taste: Estimating import price inflation across US income groups,” *Journal of International Economics*, 2020, 127, 103382.

- Husted, Lucas, John Rogers, and Bo Sun**, “Monetary Policy Uncertainty,” *Journal of Monetary Economics*, 2020, 115, 20–36.
- Irwin, Douglas A**, *Clashing Over Commerce*, University of Chicago Press, 2017.
- Jarociński, Marek and Peter Karadi**, “Deconstructing Monetary Policy Surprises—The Role of Information Shocks,” *American Economic Journal: Macroeconomics*, April 2020, 12 (2), 1–43.
- Jia, Chengcheng**, “The Informational Effect of Monetary Policy and the Case for Policy Commitment,” Working Paper 2020.
- Jordà, Òscar**, “Estimation and inference of impulse responses by local projections,” *American economic review*, 2005, 95 (1), 161–182.
- Kleibergen, Frank**, “Testing parameters in GMM without assuming that they are identified,” *Econometrica*, 2005, 73 (4), 1103–1123.
- Klein, Dan and Christopher D Manning**, “Accurate unlexicalized parsing,” in “Proceedings of the 41st annual meeting of the association for computational linguistics” 2003, pp. 423–430.
- Kuttner, Kenneth N**, “Monetary policy surprises and interest rates: Evidence from the Fed funds futures market,” *Journal of Monetary Economics*, 2001, 47 (3), 523–544.
- Lakdawala, Aeimit**, “Decomposing the effects of monetary policy using an external instruments SVAR,” *Journal of applied econometrics*, 2019, 34 (6), 934–950.
- Lepetyuk, Vadym, Lilia Maliar, Serguei Maliar, and John B Taylor**, “Policy Anticipation Effects in New Keynesian Models,” 2021.
- Lewis, Daniel J**, “Announcement-Specific Decompositions of Unconventional Monetary Policy Shocks and Their Macroeconomic Effects,” *FRB of New York Staff Report*, 2019, (891).
- Lindsey, David E**, “A Modern History of FOMC Communication: 1975-2002,” *Board of Governors of the Federal Reserve System, Washington, DC*, 2003.
- Loughran, Tim and Bill McDonald**, “When is a Liability not a Liability? Textual Analysis, Dictionaries, and 10-Ks,” *The Journal of Finance*, 2011, 66 (1), 35–65.
- Lucas, Robert E**, “Expectations and the neutrality of money,” *Journal of Economic Theory*, 1972, 4 (2), 103–124.
- Lucca, David O. and Francesco Trebbi**, ““Measuring Central Bank Communication: An Automated Approach with Application to FOMC Statements”,” *NBER Working Paper Series*, 2009, 15367.

- Lunsford, Kurt G**, “Policy language and information effects in the early days of Federal Reserve forward guidance,” *American Economic Review*, 2020, 110 (9), 2899–2934.
- Lustenbergera, Thomas and Enzo Rossib**, “Does Central Bank Transparency and Communication Affect Financial and Macroeconomic Forecasts?,” *International Journal of Central Banking*, 2020.
- Manning, Christopher D, Mihai Surdeanu, John Bauer, Jenny Rose Finkel, Steven Bethard, and David McClosky**, “The Stanford CoreNLP natural language processing toolkit,” in “Proceedings of 52nd annual meeting of the association for computational linguistics: system demonstrations” 2014, pp. 55–60.
- Marcus, Mitchell, Grace Kim, Mary Ann Marcinkiewicz, Robert MacIntyre, Ann Bies, Mark Ferguson, Karen Katz, and Britta Schasberger**, “The Penn Treebank: Annotating Predicate Argument Structure,” in “Proceedings of the workshop on Human Language Technology” Association for Computational Linguistics 1994, pp. 114–119.
- Meade, Ellen E.**, “The FOMC: Preferences, Voting, and Consensus,” *Federal Reserve Bank of St. Louis Review*, 2005, 87, 93–101.
- **and David Stasavage**, “Publicity of Debate and the Incentive to Dissent: Evidence from the US Federal Reserve,” *The Economic Journal*, 2008, 118 (528), 695–717.
- **, Nicholas A. Burk, and Melanie Josselyn**, “The FOMC Meeting Minutes: An Assessment of Counting Words and the Diversity of Views,” *Federal Reserve Board FEDS Notes*, 2015.
- Melosi, Leonardo**, “Signalling effects of monetary policy,” *The Review of Economic Studies*, 2017, 84 (2), 853–884.
- Meltzer, Allan H**, *A History of the Federal Reserve, Volume 2*, University of Chicago Press, 2010.
- Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean**, “Distributed representations of words and phrases and their compositionality,” in “Advances in neural information processing systems” 2013, pp. 3111–3119.
- Miranda-Agrippino, Silvia and Giovanni Ricco**, “The Transmission of Monetary Policy Shocks,” *American Economic Journal: Macroeconomics*, July 2021, 13 (3), 74–107.
- Morris, Stephen and Hyun Song Shin**, “Inertia of forward-looking expectations,” *American Economic Review*, 2006, 96 (2), 152–157.
- Nakamura, Emi and Jón Steinsson**, “High Frequency Identification of Monetary Non-Neutrality,” *Quarterly Journal of Economics*, 2018, 133 (3), 1283–1330.
- Naszodi, Anna, Csaba Csavas, , Szilard Erhart, and Daniel Felcser**, “Which Aspects of Cen-

- tral Bank Transparency Matter? A Comprehensive Analysis of the Effect of Transparency of Survey Forecasts,” *International Journal of Central Banking*, 2016.
- Newey, Whitney K and Kenneth D West**, “Automatic lag selection in covariance matrix estimation,” *The Review of Economic Studies*, 1994, 61 (4), 631–653.
- Nimark, Kristoffer**, “Dynamic higher order expectations,” Technical Report, Working Paper 2017.
- Petersen, K. B. and M. S. Pedersen**, “The Matrix Cookbook,” nov 2012. Version 20121115.
- Porter, Martin F**, “An algorithm for suffix stripping,” *Program*, 1980.
- Ramey, Valerie A**, “Macroeconomic shocks and their propagation,” in “Handbook of macroeconomics,” Vol. 2, Elsevier, 2016, pp. 71–162.
- **and Sarah Zubairy**, “Government spending multipliers in good times and in bad: evidence from US historical data,” *Journal of Political Economy*, 2018, 126 (2), 850–901.
- Rendahl, Pontus**, “Linear time iteration,” Technical Report, Working Paper 2017.
- Rigobon, Roberto**, “Identification through Heteroskedasticity,” *Review of Economics and Statistics*, 2003, 85 (4), 777–792.
- Riker, David**, “Geographically Disaggregated Import Data and Consumer Gains from Trade,” *J. Int’l Com. & Econ.*, 2013, 5, 65.
- Romer, Christina D. and David H. Romer**, “Federal Reserve Information and the Behavior of Interest Rates,” *American Economic Review*, 2000, 90 (3), 429–457.
- **and** — , “A New Measure of Monetary Shocks: Derivation and Implications,” *American Economic Review*, 2004, 94 (4), 1055–1084.
- **and** — , “The FOMC Versus the Staff: Where Can Monetary Policymakers Add Value?,” *American Economic Review, Paper and Proceedings*, 2008, 98 (2), 230–35.
- ROPA**, *FOMC Records of Policy Actions 1976-1993*, Board of Governors of the Federal Reserve System, 1976–1993.
- Sastry, Karthik**, “Disagreement About Monetary Policy,” Technical Report, Working Paper 2021.
- Schatz, Arthur W**, “The Anglo-American Trade Agreement and Cordell Hull’s Search for Peace 1936-1938,” *The Journal of American History*, 1970, 57 (1), 85–103.
- Schott, Peter K**, “The relative sophistication of Chinese exports,” *Economic Policy*, 2008, 23 (53),

6–49.

**Shapiro, Adam Hale, Moritz Sudhof, and Daniel J Wilson**, “Measuring news sentiment,” *Journal of Econometrics*, 2020.

**Shorrocks, Anthony F**, “Inequality decomposition by factor components,” *Econometrica: Journal of the Econometric Society*, 1982, pp. 193–211.

**Smets, Frank and Rafael Wouters**, “Shocks and frictions in US business cycles: A Bayesian DSGE approach,” *American economic review*, 2007, 97 (3), 586–606.

**Stock, James H and Mark W Watson**, “Twenty years of time series econometrics in ten pictures,” *Journal of Economic Perspectives*, 2017, 31 (2), 59–86.

**Tavares, Samia Costa**, “The political economy of the European customs classification,” *Public Choice*, 2006, 129 (1), 107–130.

**Taylor, John B.**, “Discretion versus Policy Rules in Practice,” in “Carnegie-Rochester Conference Series on Public Policy,” Vol. 39 1993, pp. 195–214.

**Tenreyro, Silvana and Gregory Thwaites**, “Pushing on a string: US monetary policy is less powerful in recessions,” *American Economic Journal: Macroeconomics*, 2016, 8 (4), 43–74.

**Thornton, Daniel L and David C Wheelock**, “Making sense of dissents: a history of FOMC dissents,” *Federal Reserve Bank of St. Louis Review*, 2014, 96 (3), 213–227.

**Townsend, Robert M**, “Forecasting the forecasts of others,” *Journal of Political Economy*, 1983, 91 (4), 546–588.

**Transcripts**, *FOMC Transcripts 1976-2008*, Board of Governors of the Federal Reserve System, 1976–2008.

**United Nations**, “Classification by Broad Economic Categories,” Statistical Paper, Statistics Division, Department of Economic and Social Affairs 2002.

**United States International Trade Commission**, *Potential impact on the U.S. economy and industries of the GATT Uruguay Round Agreements*, Vol. 1, U.S. International Trade Commission, 1994.

**United States Tariff Commission**, *Trade Agreement Between the United States and the United Kingdom* 1938.

—, *Twenty-third Annual Report of The United States Tariff Commission* 1939.

—, “Contribution of the United States to the Trade Negotiations under the Kennedy

Round 1964-67 (Part III of III),” *GATT Bilateral Negotiating Material by Round* ([https://www.wto.org/english/docs\\_e/gattbilaterals\\_e/indexbyround\\_e.htm](https://www.wto.org/english/docs_e/gattbilaterals_e/indexbyround_e.htm)), 1965.

**USITC**, “The effectiveness of escape clause relief in promoting adjustment to import competition,” *USITC publication*, 1982, 1229, 332–115.

**Woodford, Michael**, “Monetary Policy in the Information Economy,” in “Symposium on Economic Policy for the Information Economy” Federal Reserve Bank of Kansas City 2001, pp. 297–370.

– , “Imperfect Common Knowledge and the Effects of Monetary Policy,” *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, 2003, p. 25.

– , “Central bank communication and policy effectiveness,” in “The Greenspan Era: Lessons for the Future” Federal Reserve Bank of Kansas City 2005, pp. 399–474.

– , “Comment on ‘Macroeconomic effects of Federal Reserve forward guidance’,” *Brookings Papers on Economic Activity*, 2012, pp. 1–80.

**Woolley, John T and Joseph Gardner**, “The effect of “sunshine” on policy deliberation: The case of the Federal Open Market Committee,” *The Social Science Journal*, 2017, 54 (1), 13–29.

**Wright, Jonathan H.**, “What does monetary policy do to long-term interest rates at the zero lower bound?,” *The Economic Journal*, 2012, 122 (564), F447–F466.

**Wu, Jing Cynthia and Fan Dora Xia**, “Measuring the macroeconomic impact of monetary policy at the zero lower bound,” *Journal of Money, Credit and Banking*, 2016, 48 (2-3), 253–291.

**Yellen, Janet**, “Communication in monetary policy,” in “Speech to the Society of American Business Editors and Writers 50th Anniversary Conference, Washington, DC” 2013.

**Zeileis, Achim**, “Econometric Computing with HC and HAC Covariance Matrix Estimators,” *Journal of Statistical Software*, 2004, 11 (10), 1–17.

– , **Susanne Köll, and Nathaniel Graham**, “Various Versatile Variances: An Object-Oriented Implementation of Clustered Covariances in R,” *Journal of Statistical Software*, 2020, 95 (1), 1–36.

**Zhang, Xu**, “Disentangling the Information Effects in the Federal Reserve’s Monetary Policy Announcements,” Technical Report, Working Paper 2020.

## Appendix A: Numerical Example: Textbook New Keynesian Model

In this appendix I show that the illustrative model in the text is a representation of the textbook New Keynesian model found in chapter 3 of (Galí, 2015, Ch. 3). I also show that the relevant structural impact matrix in the model ( $\mathbf{M}$  in the text) has the same signs as what I estimate, and the type of solution that forms the notion of how “theoretically consistent” monetary shocks should behave. That model takes the form

$$\pi_t = c_\pi + \beta \mathbb{E}_t[\pi_{t+1}] + \kappa y_t - c_s s_{t|t} \quad (\text{A.1a})$$

$$y_t = \mathbb{E}_t[y_{t+1}] - \frac{1}{\sigma} (i_t - \mathbb{E}_t[\pi_{t+1}] - \rho) + \frac{1}{\sigma} (1 - \rho_d) d_{t|t} \quad (\text{A.1b})$$

$$i_t = \rho + \phi_\pi \pi_t + \phi_y (y_t - y) + m_{t|t} \quad (\text{A.1c})$$

where  $y_t$  is output,  $\pi_t$  is inflation, and  $i_t$  is the nominal interest rate. The model’s three structural shocks are a supply shock, demand shock, and monetary policy shock which, respectively, follow the following AR(1) processes:

$$s_t = \rho_s s_t + \varepsilon_t^s \quad (\text{Technology Shock})$$

$$d_t = \rho_d d_t + \varepsilon_t^d \quad (\text{Demand/Pref. Shock})$$

$$m_t = \rho_m m_t + \varepsilon_t^m \quad (\text{Monetary Shock})$$

where the innovation terms  $\varepsilon_t^i$  are independent and normally distributed with means of zero and standard deviation  $\sigma_s$ ,  $\sigma_d$ , and  $\sigma_m$ , respectively

**Calibration** I calibrate the model exactly as in chapter 3 of Galí (2015). The value of the structural parameters are

Parameter	Symbol	Value	Parameter	Symbol	Value
Discount factor	$\beta$	0.99	Risk aversion	$\sigma$	1
Inverse Frisch Elast.	$\varphi$	1	Cobb-Douglas	$\alpha$	0.25
Consumption elast. of subs.	$\epsilon$	9	Interest semielast. of mon. demand	$\eta$	4
Price stickiness	$\theta$	0.75	Taylor rule inflation	$\phi_\pi$	1.5
Taylor rule output	$\phi_y$	0.125	Persistence of monetary shock	$\rho_m$	0.5
Persistence of demand shock	$\rho_d$	0.5	S.D. of demand shock	$\sigma_d$	1.0
S.D. of monetary shock	$\sigma_m$	0.5	S.D. of supply shock	$\sigma_s$	0

This implies the following values for the parameters used in the model of equation (A.1):

$$\begin{aligned} \mu &= \log\left(\frac{\epsilon}{\epsilon-1}\right) \approx 0.117 \\ \Psi_y &= -\frac{(1-\alpha)(\mu - \log(1-\alpha))}{\sigma(1-\alpha) + \varphi + \alpha} \approx -0.152 \\ \Theta &= \frac{1-\alpha}{1-\alpha + \alpha\epsilon} = 0.25 \\ \lambda &= \frac{(1-\theta)(1-\beta\theta)}{\theta} \Theta \approx 0.021 \\ \kappa &= \lambda \left(\sigma + \frac{\varphi + \alpha}{1-\alpha}\right) \approx 0.057 \\ c_\pi &= -\kappa \Psi_y \approx 0.008 \\ c_s &= -\kappa \left(\frac{1+\varphi}{\sigma(1-\alpha) + \varphi + \alpha}\right) \approx -0.057 \\ \rho &= -\log(\beta) \approx 0.010 \end{aligned}$$

**Solution Approach** The model can be put into the form of equation (1.14)

$$\underbrace{\begin{bmatrix} -\kappa & 1 & 0 \\ 1 & 0 & \frac{1}{\sigma} \\ -\phi_y & -\phi_\pi & 1 \end{bmatrix}}_{\mathbf{A}} \underbrace{\begin{bmatrix} y_t \\ \pi_t \\ i_t \end{bmatrix}}_{\mathbf{y}_t} + \underbrace{\begin{bmatrix} 0 & -\beta & 0 \\ -1 & \frac{1}{\sigma} & 0 \\ 0 & 0 & 0 \end{bmatrix}}_{\mathbf{B}} \mathbb{E}_t \underbrace{\begin{bmatrix} y_{t+1} \\ \pi_{t+1} \\ i_{t+1} \end{bmatrix}}_{\mathbf{y}_{t+1}} + \underbrace{\begin{bmatrix} -c_\pi \\ \frac{\rho}{\sigma} \\ -\rho + \phi_y y \end{bmatrix}}_{\mathbf{C}} + \underbrace{\begin{bmatrix} 0 & c_s & 0 \\ -\frac{1}{\sigma}(1-\rho_d) & 0 & 0 \\ 0 & 0 & -1 \end{bmatrix}}_{\mathbf{D}} \underbrace{\begin{bmatrix} d_{t|t} \\ s_{t|t} \\ m_{t|t} \end{bmatrix}}_{\mathbf{m}_{t|t}} = \mathbf{0}$$

Given the simple format of the Taylor rule, this system can be reduced to a system in  $i_t$  and  $y_t$  by solving the Taylor rule for  $\pi_t$ . The new system takes the form

$$\underbrace{\begin{bmatrix} -\left(\frac{\phi_y}{\phi_\pi} + \kappa\right) & \frac{1}{\phi_\pi} \\ 1 & \frac{1}{\sigma} \end{bmatrix}}_{\mathbf{A}} \begin{bmatrix} y_t \\ i_t \end{bmatrix} + \underbrace{\begin{bmatrix} \beta \frac{\phi_y}{\phi_\pi} & -\frac{\beta}{\phi_\pi} \\ \frac{\phi_y}{\phi_\pi \sigma} - 1 & -\frac{1}{\phi_\pi \sigma} \end{bmatrix}}_{\mathbf{B}} \mathbb{E}_t \begin{bmatrix} y_{t+1} \\ i_{t+1} \end{bmatrix} \\ + \underbrace{\begin{bmatrix} -\frac{\rho}{\phi_\pi} - c_\pi - \beta \frac{\phi_y}{\phi_\pi} y + \frac{\rho \beta}{\phi_\pi} \\ -\frac{\phi_y}{\phi_\pi \sigma} y + \frac{\rho}{\phi_\pi \sigma} - \frac{\rho}{\sigma} \end{bmatrix}}_{\bar{\mathbf{C}}} + \underbrace{\begin{bmatrix} 0 & c_s & \frac{\beta \rho m}{\phi_\pi} - \frac{1}{\phi_\pi} \\ -\frac{1}{\sigma}(1 - \rho d) & 0 & \frac{\rho}{\phi_\pi \sigma} \end{bmatrix}}_{\mathbf{D}} \begin{bmatrix} d_{t|t} \\ s_{t|t} \\ m_{t|t} \end{bmatrix} = \mathbf{0}$$

**Full Information Numerical Solution** Following the solution method described in section C, the solution of the model without noise consists of the exogenous processes for the demand and monetary shocks and the analogue to equation (C.4a):

$$\begin{bmatrix} y_t \\ i_t \end{bmatrix} = \begin{bmatrix} -0.15 \\ 0.01 \end{bmatrix} + \underbrace{\begin{bmatrix} 0.67 & -1.35 \\ 0.20 & 0.60 \end{bmatrix}}_{\mathbf{M}_{\text{NK}}} \begin{bmatrix} d_{t|t} \\ m_{t|t} \end{bmatrix}, \quad (\text{A.2})$$

where, given the full-information assumption,  $d_{t|t} = d_t$  and  $m_{t|t} = m_t$ ; in words, actual and perceived shocks coincide. Pre-multiplying both sides by  $\mathbf{M}_{\text{NK}}^{-1}$  and then normalizing the coefficients on  $y_t$  and  $i_t$  to unity yields the reduced-form DIS curve and Taylor rules from this model.

$$y_t = -0.17 - 2.26i_t + 1.13d_t \\ i_t = 0.01 + 0.29y_t + \varepsilon_t.$$

## Appendix B: Formal Signaling Structures

In the text I remain agnostic as the nature of the information received by Fed watchers. In this appendix I provide simple formal examples that generate the intuitions described in the text.

### B.1 Static Signaling Model

Assume that, for a Fed announcement occurring at time  $t$ , the Fed emits noisy signals about each shock given by

$$\begin{aligned} s_t^\varepsilon &= \varepsilon_\tau + n_t^\varepsilon \\ s_t^\eta &= \eta_\tau + n_t^\eta. \end{aligned}$$

The noise components  $n_t^\varepsilon$  and  $n_t^\eta$  are independently distributed normal variables with zero mean and respective variances  $\sigma_{n,\varepsilon}^2$  and  $\sigma_{n,\eta}^2$ . The formulation of these signals suggests that the Fed chooses “how noisy” to make its signals. This noise could arise, instead, from the noise with which the Fed perceives the underlying signals itself.

The signals  $s_t^\varepsilon$  and  $s_t^\eta$  are used to form expectations of the economy’s structural shocks using Bayes’ rule and knowledge of the model’s parameters. Prior beliefs for each of the structural shocks are zero (the mean of the underlying distributions), so posterior expectations—equivalently, expectation revisions—are

$$\mathbb{E}_t[\varepsilon_\tau] = \underbrace{\left( \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_{n,\varepsilon}^2} \right)}_{\kappa_\varepsilon} s_t^\varepsilon \qquad \mathbb{E}_t[\eta_\tau] = \underbrace{\left( \frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_{n,\eta}^2} \right)}_{\kappa_\eta} s_t^\eta. \qquad (\text{B.1})$$

The simple form of these expectations follows from the information structure above. That be-

ing said, I made two simplifications in formulating the information structure. The first concerns the lack of dynamics in the structural shocks ( $\varepsilon_\tau$  and  $\eta_\tau$ ). With dynamic shocks, nearly identical formulas could replace equation (B.1) by positing that the signals were of the white-noise innovation to the structural shocks (or a lag-polynomial thereof, if  $\tau > t$ ) instead of the structural shocks themselves. However, because the expectations revisions that I study are empirically uncorrelated, this simplification is warranted. The second simplification concerns the nature of the signals—specifically that the noise is independent across the structural shocks. If, instead, the signals were of the observable variables ( $y_\tau$  and  $i_\tau$ ), then the posterior expectations of the structural shocks would be correlated, as Acosta and Afrouzi (In Progress) demonstrate. Again, my empirical results do not suggest that forecast revisions regarding structural shocks are correlated, which alleviates this concern.<sup>1</sup>

## B.2 Dynamic Extension

The discussion in the previous section was ambiguous with respect to the timing of the outcome variables. Empirically, the measures of output and interest rate expectations both captured the average value of those variables over the coming year. With measures of expectations now differentiated by time, this ambiguity must be resolved.

To do this, I assume that the monetary and information shocks follow exogenous first-order autoregressive processes, given by

$$\begin{aligned} \varepsilon_t &= \rho_\varepsilon \varepsilon_{t-1} + \mu_t^\varepsilon & \mu_t^\varepsilon &\sim \mathcal{N}\left(0, \sigma_\varepsilon^2\right) \\ \eta_t &= \rho_\eta \eta_{t-1} + \mu_t^\eta & \mu_t^\eta &\sim \mathcal{N}\left(0, \sigma_\eta^2\right). \end{aligned}$$

I have simplified the timing conventions substantially here, assuming that  $t$  corresponds to calendar

---

<sup>1</sup>This assumption seems less plausible if markets *only* observed the interest-rate decision itself. In reality, Fed announcements are composed of a post-meeting statement over my entire sample that explains the rationale behind the policy decision. Thus, the announcement is multidimensional. What's more, since the start of my sample several additional dimensions have been added to this signal: press conferences (with the opportunity to answer questions from the press) and economic forecasts.

time and Fed meetings, and agents only receive information from the Fed. These simplifications are not necessary but help me to make my argument concisely. The next step is to allow the Fed to send a signal about future interest rates. For expositional purposes I restrict the model to allow the Fed to emit a signal only about next-period's interest rate. The information set of agents thus consists of all prior signals, along with

$$\begin{aligned}
 f_t^\varepsilon &= \mu_{t+1}^\varepsilon + n_t^f & n_t^f &\sim \mathcal{N}(0, \sigma_f^2) \\
 s_t^\varepsilon &= \mu_t^\varepsilon + n_t^\varepsilon & n_t^\varepsilon &\sim \mathcal{N}(0, \sigma_{n,\varepsilon}^2) \\
 s_t^\eta &= \mu_t^\eta + n_t^\eta & n_t^\eta &\sim \mathcal{N}(0, \sigma_{n,\eta}^2).
 \end{aligned}$$

With this information structure, the variance of forecast revisions of monetary policy innovations are given by

$$\text{var}(\hat{\mu}_t) = \frac{\sigma_\varepsilon^4 \sigma_f^4}{(\sigma_\varepsilon^2 + \sigma_f^2)(\sigma_\varepsilon^2 \sigma_{n,\varepsilon}^2 + \sigma_\varepsilon^2 \sigma_f^2 + \sigma_{n,\varepsilon}^2 \sigma_f^2)} \quad \text{var}(\hat{\mu}_{t+1}) = \frac{\sigma_\varepsilon^4}{\sigma_\varepsilon^2 + \sigma_f^2},$$

where now the hat notation denotes changes between periods  $t$  and  $t - 1$ , which is equivalent to high-frequency changes in this simple timing setup:  $\hat{x}_t = \mathbb{E}_t[x_t] - \mathbb{E}_{t-1}[x_t]$ . These equations underlie the intuition for how the introduction of forward guidance—here a decrease in  $\sigma_f$ —can simultaneously make current-meeting interest-rate surprises smaller, but future-period interest-rate surprises larger. Signaling about tomorrow's monetary shock has two effects. Today, it allows markets to put more weight on that signal when forming forecasts about tomorrow's innovation, thus increasing the average size of forecast revisions (formally,  $\text{var}(\hat{\mu}_{t+1})$  is decreasing in  $\sigma_f$ ). When tomorrow,  $t + 1$ , rolls around, agents will have already received a signal in period  $t$ , so they have less prior uncertainty about the current-period innovation. The signal is relatively less informative, dampening the size of tomorrow's revisions ( $\text{var}(\hat{\mu}_t)$  is increasing in  $\sigma_f$ ). Note that, in the terminology of Gürkaynak et al. (2005),  $\hat{\mu}_t$  and  $\hat{\mu}_{t+1}$  are akin to target and path shocks, respectively—the former affects both  $\hat{i}_t$  and  $\hat{i}_{t+1}$ , while the latter does not affect  $\hat{i}_t$ .

This intuition continues to operate when considering the size of forecast revisions to the forward rates  $i_t$  and  $i_{t+1}$ . Combining equation (1.8a) with the fact that  $\hat{\varepsilon}_{t+1} = \hat{\mu}_{t+1} + \rho\hat{\mu}_t$  (and that  $\mathbb{E}[\hat{\mu}_{t+1}\hat{\mu}_t] = 0$ ), it can be shown that when  $\rho \leq 1$ ,

$$\frac{\partial \text{var}(\hat{i}_t)}{\partial \sigma_f} > 0 \qquad \frac{\partial \text{var}(\hat{i}_{t+1})}{\partial \sigma_f} < 0. \qquad (\text{B.2})$$

In words, a clearer signal about future interest rates—i.e., a decrease in  $\sigma_f$ —leads to smaller current-period interest-rate revisions, but larger future interest-rate revisions. Note, however, that  $\hat{i}_{t+1}$  is a forward rate. Under the expectations theory of the term structure, the theoretical analog to my empirical variable is expected yield on a 2-period bond—under the expectations theory of the term structure, this is

$$i_{t,1} \equiv \frac{1}{2} \mathbb{E}_t[i_t + i_{t+1}].$$

Thus while the model offers clear predictions for how the size of forecast revisions for forward rates responds to the introduction of forward guidance, equation (B.2) shows that the response of the size of yield revisions is ambiguous. In particular, the response depends on the persistence of the monetary shock,  $\varepsilon_t$ . As the shock becomes increasingly transitory, yield forecast revisions  $\hat{i}_{t,1}$  become unambiguously larger in response to forward guidance, since current innovations carry through less to future shocks.<sup>2</sup> My finding that yield revisions increase following the introduction of forward guidance suggests, then, that monetary shocks are indeed fairly transitory.

The dynamic form of the baseline model also provides an explanation for why longer-term interest rates are more susceptible to information effects. To see this, note that longer-term interest-rate (forward) surprises are

$$\hat{i}_{t+k} = \rho_\eta^k \eta_{t|t} + \rho_\varepsilon^{k-1} \varepsilon_{t+1|t}.$$

This can be used to express the contribution of information effects—relative to the contribution of

---

<sup>2</sup>This result is straightforward: If shocks aren't persistent, then today's interest-rate shock does not affect tomorrow's interest-rate forward.

monetary shocks—to the variance of interest rate surprises as

$$\lambda_k \equiv \frac{\rho_\eta^{2k} \text{var}(\hat{\eta}_t) / \text{var}(\hat{i}_{t+k})}{\rho_\varepsilon^{2(k-1)} \text{var}(\hat{\varepsilon}_t) / \text{var}(\hat{i}_{t+k})}.$$

Thus, as the interest-rate horizon increases, the relative contribution of information also increases *if* the information shock is more persistent, since

$$\frac{d \log(\lambda_k)}{dk} = 2(\log(\rho_\eta) - \log(\rho_\varepsilon)) > 0 \iff \rho_\eta > \rho_\varepsilon.$$

Empirically, Smets and Wouters (2007)—who estimate a DSGE of the US economy driven by a rich set of macroeconomic shocks—suggest that monetary shocks are fairly transitory.

## Appendix C: Solving the Noisy Information Model

This appendix contains the instructions to solve a noisy-information linear dynamic model following Blanchard et al. (2013), whose exposition I follow closely. The discussion is useful for considering models in which current decisions depend not only expected future decisions, but also on lagged decisions (an extension of the model in the text). The model takes the form

$$\mathbf{A} \mathbf{x}_t + \mathbf{B} \mathbb{E}_t[\mathbf{x}_{t+1}] + \mathbf{C} \mathbf{x}_{t-1} + \bar{\mathbf{C}} + \mathbf{D} \boldsymbol{\xi}_{t|t} = \mathbf{0} \quad (\text{C.1})$$

where  $\mathbf{x}_t$  is a vector of observable macroeconomic variables,  $z_{t|\tau}$  denotes the mathematical expectation of  $z_t$  given information at time  $\tau$ , and  $\boldsymbol{\xi}_t$  is a vector of mutually independent structural shocks that evolves according to:

$$\boldsymbol{\xi}_t = \mathbf{H} \boldsymbol{\xi}_{t-1} + \mathbf{J} \boldsymbol{\varepsilon}_t \quad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_\varepsilon), \quad (\text{C.2})$$

where  $\boldsymbol{\Sigma}_\varepsilon$  is a diagonal matrix. At the beginning of each period  $t$  but before making decisions about  $\mathbf{x}_t$ , agents receive a noisy signal of the structural shock of the form

$$\mathbf{s}_t = \mathbf{F} \boldsymbol{\xi}_t + \mathbf{G} \boldsymbol{\nu}_t \quad \boldsymbol{\nu}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_\nu), \quad (\text{C.3})$$

which they use to form expectations about  $\boldsymbol{\xi}_t$  using the Kalman filter. In the full-information case (i.e.  $\boldsymbol{\Sigma}_\nu \rightarrow \mathbf{0}$ ), instead, structural shocks are observed perfectly and so  $\boldsymbol{\xi}_{t|t} = \boldsymbol{\xi}_t$ . I conjecture that the model satisfies all necessary stability conditions such that it admits a solution takes the

following form

$$\mathbf{x}_t = \bar{\mathbf{x}} + \mathbf{L} \mathbf{x}_{t-1} + \mathbf{M} \xi_{t|t} \quad (\text{C.4a})$$

$$\xi_{t|t} = (\mathbf{I} - \mathbf{K}\mathbf{F})\mathbf{H}\xi_{t-1|t-1} + \mathbf{K}\mathbf{s}_t \quad (\text{C.4b})$$

Plugging the conjectured guess of equation (C.4a) into (C.1) gives

$$\mathbf{A}(\bar{\mathbf{x}} + \mathbf{L} \mathbf{x}_{t-1} + \mathbf{M} \xi_{t|t}) + \mathbf{B} \mathbb{E}_t[\bar{\mathbf{x}} + \mathbf{L} \mathbf{x}_t + \mathbf{M} \xi_{t+1|t+1}] + \mathbf{C} \mathbf{x}_{t-1} + \bar{\mathbf{C}}\mathbf{D}\xi_{t|t} = \mathbf{0}. \quad (\text{C.5})$$

Plugging in the guess again for  $\mathbf{x}_t$  and noticing that, using the law of iterated expectations and (C.2),  $\mathbb{E}_t[\xi_{t+1|t+1}] = \xi_{t+1|t} = \mathbf{H}\xi_{t|t}$ , equation (C.5) becomes

$$[(\mathbf{A} + \mathbf{B}(\mathbf{I} + \mathbf{L}))\bar{\mathbf{x}} + \bar{\mathbf{C}}] + [\mathbf{A}\mathbf{L} + \mathbf{B}\mathbf{L}^2 + \mathbf{C}] \mathbf{x}_{t-1} + [\mathbf{A}\mathbf{M} + \mathbf{B}\mathbf{L}\mathbf{M} + \mathbf{B}\mathbf{M}\mathbf{H} + \mathbf{D}] \xi_{t|t} = \mathbf{0}.$$

To ensure that this equation holds with equality regardless of  $\mathbf{x}_{t-1|t-1}$  and  $\xi_{t|t}$ , it must be then that

$$\mathbf{A}\mathbf{L} + \mathbf{B}\mathbf{L}^2 + \mathbf{C} = \mathbf{0} \quad \mathbf{A}\mathbf{M} + \mathbf{B}\mathbf{L}\mathbf{M} + \mathbf{B}\mathbf{M}\mathbf{H} + \mathbf{D} = \mathbf{0} \quad (\mathbf{A} + \mathbf{B}(\mathbf{I} + \mathbf{L}))\bar{\mathbf{x}} + \bar{\mathbf{C}} = \mathbf{0}$$

The first equation can be used to solve for  $\mathbf{L}$  using the method of Rendahl (2017), the second is an “encapsulating sum” problem whose solution is given in Petersen and Pedersen (2012), and the last is linear in  $\bar{\mathbf{x}}$ . Solving for Kalman gain matrix  $\mathbf{K}$  can also be done by iteration. Define the initial guess for the matrix  $\mathbf{P}$  as  $\mathbf{P}_0$ . Then iterate on the following equations over  $i$  until convergence:

$$\mathbf{K}_i = (\mathbf{H}\mathbf{P}_{i-1}\mathbf{H}' + \mathbf{J}\boldsymbol{\Sigma}_\varepsilon\mathbf{J}')\mathbf{F}(\mathbf{F}(\mathbf{H}\mathbf{P}_{i-1}\mathbf{H}' + \mathbf{J}\boldsymbol{\Sigma}_\varepsilon\mathbf{J}')\mathbf{F}' + \mathbf{G}\boldsymbol{\Sigma}_\nu\mathbf{G}')^{-1} \quad (\text{C.6a})$$

$$\mathbf{P}_i = (\mathbf{I} - \mathbf{K}_i\mathbf{F})(\mathbf{H}\mathbf{P}_{i-1}\mathbf{H}' + \mathbf{J}\boldsymbol{\Sigma}_\varepsilon\mathbf{J}') \quad (\text{C.6b})$$

Having shown that equation (C.4a) can actually solve the model, I consider how we can interpret my formal empirical model in the context of this more elaborate model. The connection is as follow: let  $\mathbf{z}_t \equiv \mathbf{x}_t - \mathbf{L} \mathbf{x}_{t-1}$  be the reduced-form residual of  $\mathbf{x}_t$ : the unexpected change vis-à-vis

the previous period's information. Then (C.4a) reveals that these reduced-form residuals, or data *surprises*, are related by a constant linear function to forecast revisions about structural shocks:

$$\mathbf{z}_{t|t} - \mathbf{z}_{t|t-1} = \mathbf{M} (\boldsymbol{\xi}_{t|t} - \boldsymbol{\xi}_{t|t-1}).$$

Thus, under this more-elaborate model I interpret my high-frequency forecast revisions as revisions about recent news about each variable, rather than revisions about the *levels* of the variables themselves.

## Appendix D: Why is $\hat{y}_t$ an invalid instrument (and bad control)?

The illustrative model showed why  $\hat{y}_t$  is an invalid instrument for information effects—it fails the exclusion restriction by failing to be independent of the monetary shock. To see that it is also a “bad control” (in the language of Angrist and Pischke (2008)), consider what happens when controlling for  $\hat{y}_t$ —equivalently (Frisch and Waugh, 1933), attempting to identify a monetary policy shocks by “purging” the monetary surprise  $\hat{i}_t$  of its information content ( $\eta_t$ ) by orthogonalizing  $\hat{i}_t$  to  $\hat{y}_t$ . Proposition 1 states the conditions under which this procedure identifies a monetary policy shock.

**Proposition 1** *Denote by  $r_t$  the residual from a linear projection of  $\hat{i}_t$  on  $\hat{y}_t$ . Unless  $r_t = 0 \ \forall \ t$ , then  $r_t$  is independent of information effects if and only if output expectations do not respond to monetary policy shocks, i.e.,  $\omega_\varepsilon = 0$ .*

Let

$$\hat{i}_t = \phi_\varepsilon \hat{\varepsilon}_t + \phi_\eta \hat{\eta}_t \tag{D.1}$$

$$\hat{y}_t = \omega_\varepsilon \hat{\varepsilon}_t + \omega_\eta \hat{\eta}_t. \tag{D.2}$$

The residual  $r_t$  (i.e., the “clean” monetary shock) is

$$r_t = \hat{i}_t - \hat{\beta} y_t$$

where

$$\hat{\beta} = \frac{\text{Cov}(\hat{y}_t, \hat{i}_t)}{\text{var}(y_t)} = \frac{\phi_\eta \omega_\eta \sigma_\eta^2 + \phi_\varepsilon \omega_\varepsilon \sigma_\varepsilon^2}{\omega_\eta^2 \sigma_y^2 + \omega_\varepsilon^2 \sigma_\varepsilon^2}$$

This residual is then a linear combination of the monetary and information shocks:

$$\begin{aligned}
r_t &= \hat{i}_t - \hat{\beta} \hat{y}_t \\
&= \phi_\varepsilon \varepsilon_t + \phi_\eta \eta_t - \hat{\beta} (\omega_\varepsilon \varepsilon_t + \omega_\eta \eta_t) \\
&= \underbrace{(\phi_\eta - \hat{\beta} \omega_\eta)}_{c_\eta} \eta_t + \underbrace{(\phi_\varepsilon - \hat{\beta} \omega_\varepsilon)}_{c_\varepsilon} \varepsilon_t.
\end{aligned}$$

This strategy then only provides a “clean” shock if  $c_\eta \equiv \phi_\eta - \hat{\beta} \omega_\eta = 0$ .

$$\begin{aligned}
0 &= \phi_\eta - \hat{\beta} \omega_\eta \\
\iff 0 &= \phi_\eta - \omega_\eta \frac{\phi_\eta \omega_\eta \sigma_\eta^2 + \phi_\varepsilon \omega_\varepsilon \sigma_\varepsilon^2}{\omega_\eta^2 \sigma_\eta^2 + \omega_\varepsilon^2 \sigma_\varepsilon^2} \\
\iff 0 &= \phi_\eta (\omega_\eta^2 \sigma_\eta^2 + \omega_\varepsilon^2 \sigma_\varepsilon^2) - \omega_\eta (\phi_\eta \omega_\eta \sigma_\eta^2 + \phi_\varepsilon \omega_\varepsilon \sigma_\varepsilon^2) \\
\iff 0 &= \phi_\eta \omega_\eta^2 \sigma_\eta^2 + \phi_\eta \omega_\varepsilon^2 \sigma_\varepsilon^2 - \phi_\eta \omega_\eta^2 \sigma_\eta^2 - \omega_\eta \phi_\varepsilon \omega_\varepsilon \sigma_\varepsilon^2 \\
\iff 0 &= \phi_\eta \omega_\varepsilon^2 \sigma_\varepsilon^2 - \omega_\eta \phi_\varepsilon \omega_\varepsilon \sigma_\varepsilon^2 \\
\iff 0 &= \omega_\varepsilon \sigma_\varepsilon^2 (\phi_\eta \omega_\varepsilon - \omega_\eta \phi_\varepsilon)
\end{aligned}$$

The strategy thus provides clean shock in three cases.

1. First, the case where  $\sigma_\varepsilon^2 = 0$  means that that there are no monetary shocks, so  $r_t = 0 \quad \forall \quad t$ , violating our assumptions.
2. The knife-edge case with  $\phi_\eta \omega_\varepsilon = \omega_\eta \phi_\varepsilon$  also results in  $r_t = 0 \quad \forall \quad t$ . To see this, note that this

assumption also implies that  $c_\varepsilon = 0$ :

$$\begin{aligned}
c_\varepsilon &= \phi_\varepsilon - \hat{\beta}\omega_\varepsilon \\
&= \phi_\varepsilon - \omega_\varepsilon \frac{\phi_\eta \omega_\eta \sigma_\eta^2 + \phi_\varepsilon \omega_\varepsilon \sigma_\varepsilon^2}{\omega_\eta^2 \sigma_\eta^2 + \omega_\varepsilon^2 \sigma_\varepsilon^2} \\
&= \frac{\phi_\varepsilon (\omega_\eta^2 \sigma_\eta^2 + \omega_\varepsilon^2 \sigma_\varepsilon^2) - \omega_\varepsilon \phi_\eta \omega_\eta \sigma_\eta^2 - \phi_\varepsilon \omega_\varepsilon^2 \sigma_\varepsilon^2}{\omega_\eta^2 \sigma_\eta^2 + \omega_\varepsilon^2 \sigma_\varepsilon^2} \\
&= \frac{\phi_\varepsilon \omega_\eta^2 \sigma_\eta^2 - \omega_\varepsilon \phi_\eta \omega_\eta \sigma_\eta^2}{\omega_\eta^2 \sigma_\eta^2 + \omega_\varepsilon^2 \sigma_\varepsilon^2} \\
&= \frac{\omega_\eta \sigma_\eta^2 (\omega_\eta \phi_\varepsilon - \omega_\varepsilon \phi_\eta)}{\omega_\eta^2 \sigma_\eta^2 + \omega_\varepsilon^2 \sigma_\varepsilon^2} \\
&= 0
\end{aligned}$$

Thus, with  $c_\eta = c_\varepsilon = 0$ , we have  $r_t = 0 \quad \forall t$ .

3. The final possibility is that output expectations do not respond to monetary policy shocks, i.e.,  $\omega_\varepsilon = 0$ .

## Appendix E: Text Analysis Appendix

### E.1 Word Lists

**Increase words** abound absorb absorpt acceler access accru accumul adjunct advanc ampli amplifi append arisen augment becam becom bloom blossom bolster boom boost boundless bounti branch broaden build capit collect comeback cultiv decor deepen develop doubl elabor embellish empow empower enhanc enrich exce excel expand expans extend flourish fortifi further garnish gener grow grown growth heap heighten hoard improv increas inflat intensifi inund lucr magnifi matur maxim momentum nourish overflow overwhelm peopl piec pile prolif promot prosper quicken radiat rais renaiss rise run shoot spread strengthen supplement surg sweeten thrive weight widen

**Decrease Words** abat allevi amput atrophu cheapen collaps contract corrod corros counteract cut decay declin decompos decreas deplet deprecu detract dim diminish discount discourag dispel dispens drain dwindl eas empti engulf erad eras erod eros exasper exhaust extermin fade fall falter insuffici languish leakag lighten lower melt minim pass purifi ration reced recess reduc reduct refin retard revers rid rot scarciti shrank shred shrink shrivel shrunk slow subsid subtract sunder tatter vanish wane weaken wilt wither worsen

**Rise Words** aris aros ascend ascent blast climb come elev flew fli float flood jump leap outreach peak rais rise rose scale soar stretch surfac well

**Fall Words** bag buri burst cave collaps crash descend dip dive doubl drop fall fell knock lower parachut plung rain sank set sink slid slide slip slump spray sprinkl stagger stumbl submerg sunk sunken swoon toppl torrent trip tumbl

**List of cities** BASEL BEIJING, BRASILIA, FRANKFURT, HONG KONG, JAKARTA, LONDON, MEXICO CITY, MOSCOW, MUMBAI, MUMBAI, NEW DELHI, OTTAWA, PRETORIA, RIO DE JANIERO, SAO PAULO, SINGAPORE, SYDNEY, TOKYO, TORONTO, WELLINGTON, ZURICH, KUALA LUMPUR,

## E.2 Word List Construction

I start with a set of three “seed” words: `output`, `growth`, and `economy`. I train the word2vec algorithm of Mikolov et al. (2013) on a subset of a large corpus of newspaper articles: the *The New York Times Annotated Corpus* from the University of Pennsylvania’s Linguistic Data Consortium. The full corpus contains 1.8 million articles from the New York Times between 1987 and 2007, each manually tagged by library scientists. The word2vec algorithm consists of constructing vector representations of words that, via a neural network, can predict a word in a set of text given the surrounding words. The algorithm is thus well suited to finding synonyms, hence its employment here. I trained the algorithm on 94,601 articles that were tagged as related to either economic output, prices, or labor markets.<sup>1</sup> With a vector representation of every word in the corpus, I sort words based on their distance to the average vector of my seed words.<sup>2</sup> The resulting list, along with the distances from the seed vector, is listed in table table E.1.

## E.3 Machine Learning Approach

While transparent, the approach for constructing high-frequency GDP expectations in the paper required that several, potentially subjective, choices be made. In this appendix, I pursue a machine-learning based approach to overcome this potential criticism. Specifically, I take the following

---

<sup>1</sup>Specifically, I retained articles labeled economic conditions and trends, united states economy, prices, wages and salaries, layoffs and job reductions, production, or labor. Later in the paper I also construct inflation and unemployment expectations in high frequency: thus the prices and labor tags.

<sup>2</sup>Mikolov et al. (2013) highlight that summing and distracting the vector representations of each word results in meaningful vectors. For example, the authors find that subtracting the vector for `man` from the vector for `king` results in a vector that is very similar to the vector for `queen`. To compute the distance of two words with vectors  $x$  and  $y$ , I compute the cosine similarity between them:  $x'y/[\sqrt{(x'x)} + \sqrt{(y'y)}]$ .

Table E.1: Extended GDP Word List


Term	Similarity to output + growth + economi
economic_growth	0.884380348
growth	0.844084993
economi	0.838640034
consumer_spending	0.768225532
output	0.73070604
recoveri	0.72648164
consumer_confidence	0.658283754
living_standards	0.642701291
anem	0.630283417
inflat	0.627934909

steps to measure high-frequency GDP expectations.

- First, I collected the subset of the *New York Times Annotated Corpus* described in the previous appendix. This corpus formed my training set.
- To construct features from the text, I collect my list of GDP words (which I expand to include the first 7 words from table E.1, since after 7 the words are less-obviously related to GDP). Call this  $\mathcal{Y}^+$ . I also collect the rising, falling, increasing, decreasing, strong, and weak word lists from the Harvard IV-4 dictionary. I collect the positive and negative word lists of Loughran and McDonald (2011). Call each of these lists  $\ell_i$ , and their collection  $\mathcal{L}$ . Finally, I tag all words by their tense (past, present, future) using the grammatical sentence parser described earlier.
- I concatenate all articles written between the 8<sup>th</sup> day of month  $t - 1$  and the 7<sup>th</sup> day of month  $t$ . Features of these articles (described in the next bullet) are merged with Blue Chip expectations made in month  $t$ , which are typically made during the first week of month  $t$ .
- Within each month’s concatenated list, I count the sum of all occurrences of words in each of the  $\ell_i$  word list ( $|\mathcal{L}|$  counts), all occurrences of words from the  $\ell_i$  word lists that occur within 5 words of a word from  $\mathcal{Y}^+$  ( $|\mathcal{L}| - 1$  counts here for each word from  $\mathcal{Y}^+$ ), and all raw

occurrences of words from  $\mathcal{Y}^+$ . When counting co-occurrences of  $\mathcal{L}$  words and  $\mathcal{Y}^+$  words, if `n't` or `not` occurs within the window, I flip the “sign” of the count (i.e., fall words become rise words). I count raw and co-occurrences (with words from  $\mathcal{Y}^+$ ) of tense words. I also count three-way co-occurrences of tense  $\times$   $\mathcal{L}$  words  $\times$   $\mathcal{Y}^+$  words. All counts are normalized by the number of total sentences in a month. All in all, this leaves 263 features.

- I then use a LASSO regression to estimate the mapping from these features to the *level* of Blue Chip GDP expectations (the same summary statistic from the text: average forecasts over the next year). The penalty is chosen by 10-fold cross-validation. This figure shows the model fit:



The top 3 features that positively contribute to GDP expectations are

1. “Weak” mentions of “recovery,” present tense
2. “Increasing” mentions of “output,” present tense
3. “Strong” mentions of “consumer spending,” any tense

and the top 3 features that contribute negatively are

1. “Decreasing” mention of “economy,” present tense
2. “Decreasing” word counts


3. “Weak” mention of “growth,” future tense.

- I calculate the features that I calculated from the training set on all pre-meeting and post-meeting FOMC articles from my baseline Factiva dataset. I then apply the LASSO mapping to estimate implied GDP expectation for pre- and post-meeting articles— $\omega_{PRE_t}$  and  $\omega_{POST_t}$ , respectively.
- Proceed the same way as the rest of the paper.

The structural impact matrix I estimate is

	Eurodollar	GDP
Monetary Path Shock	-0.38	0.86
Information Shock	1.18	0.68

where the variance in the second regime (relative to the first) of the monetary and path shocks are 1.03 and 0.69, respectively. Here are the impulse responses to the identified contractionary monetary path shock:



and the expansionary information shock





Figure F.1: Correlation Matrix of Variables Responding to Fed Announcements



This figure shows correlation coefficients between the variables listed and described in the text. The sample is the same as the baseline in the paper: 1999–2019, excluding the peak of the Great Recession. The Billion Prices index is available from July 2009 through July 2015. For market liquidity purposes, the breakeven inflation sample starts in 2004. The Blue Chip correlations exclude months in which the FOMC meeting occurs during the first week of the month.

## Appendix F: Alternative “Second Variables”

In principle, any variable that responds to the perceived structural shocks differently than interest rate futures could potentially be a suitable “second variable” to recover estimates of the

shocks. In this appendix, I compare how my newspaper-based measures relate to other plausible alternatives. Those comparisons are made using the correlation coefficients contained in figure F.1. Specifically, I present the correlations between the two identified shocks (information and monetary, both expansionary); the two-day change in a newspaper-based sentiment index (Shapiro et al., 2020); the two day change in the daily “Billion Prices Project” price index (Cavallo and Rigobon, 2016); <sup>1</sup> the one-day change in the 2-year breakeven inflation rate (Gürkaynak et al., 2007, 2010); the one-month changes in CPI inflation and real GDP growth expectations from the Blue Chip (described in the text); high-frequency changes in the S&P 500 (Acosta and Saia, In progress) and 4-quarter Eurodollar future; and my newspaper-based measures of high-frequency inflation and GDP forecast revisions. A few patterns are worth highlighting.

Before beginning the discussion, it is worth reiterating the notion that the “gold standard” data here would be a high-frequency measure of GDP (or inflation) forecast revisions. Why? First, this is the puzzle to solve: GDP forecast revisions present the strongest evidence of information effects. Second, aggregate output—a feature of effectively *every* macroeconomic model—is a variable over which standard theory offers strong predictions regarding expected responses to such fundamental shocks as monetary policy and aggregate demand shocks (at least in terms of signs). As such, it provides the first test for any newly-proposed measures of policy shocks. Lacking access to such a measure, particular emphasis should be placed on the correlation between these proposed alternatives and standard measures of GDP and inflation expectations, given by Blue Chip forecast revisions.

First, the newspaper-based measures of GDP and inflation forecast revisions exhibit the expected correlations with every variable under consideration: S&P returns, newspaper sentiment, price changes, and Blue Chip forecast revisions. This primarily serves as a validation exercise for these variables.

Second, and more to the point, the newspaper-based measures exhibit the strongest correlations

---

<sup>1</sup>The “two-day changes” mentioned here are, for an FOMC meeting on day  $t$  and a variable  $x$  are  $x_{t+1} - x_{t-1}$ . This is particularly important for the newspaper-based measure, so that post-announcement print articles are included. The “one-day change” is  $x_t - x_{t-1}$ .

with Blue Chip forecast revisions. In contrast, high frequency stock returns exhibit essentially no correlation with Blue Chip revisions. The next-closest alternative (aside from interest rates) is the newspaper-based sentiment index, though that correlation is much weaker—likely reflecting the fact that the index is not specifically targeted to measure expectations about a particular variable.

Finally, the market-based measures fair particularly poorly in terms of capturing changes in GDP and inflation expectations. While S&P returns exhibit the “expected” correlation with the variables under consideration, the correlation with Blue Chip forecast revisions is notably weaker.<sup>2</sup> Breakeven inflation fares even more poorly, uncorrelated with essentially every variable under consideration.<sup>3</sup>

---

<sup>2</sup>Recent evidence (Boguth et al., 2022) suggest that equity returns around FOMC announcements “largely reflect noise” and features of market microstructure.

<sup>3</sup>This suggests that a poor outlook for the application of the methodology of Andrade and Ferroni (2021) to U.S. data.

## Appendix G: Estimation and Bootstrap Procedure

**Setup and Definitions** My identification procedure begins by positing a model relating  $n$  observable variables in the vector  $x_t$  to  $n$  exogenous disturbances  $\eta_t$  given by

$$x_t = A\eta_t \qquad \eta_t \sim N(0, \Sigma_t),$$

where  $\Sigma_t$  is the (diagonal) covariance matrix of  $\eta_t$ . The time subscript on  $\Sigma_t$  indicates that the errors are heteroskedastic. Specifically, I posit that  $\Sigma_t$  can take two possible values depending on whether  $t$  falls in one of two disjoint regimes (formally, subsets  $\mathcal{R}_1$  and  $\mathcal{R}_2$ ) of my sample period. These values are given by

$$\Sigma_1 \equiv \mathbb{E}[\eta_t \eta_t' \mid t \in \mathcal{R}_1]$$

$$\Sigma_2 \equiv \mathbb{E}[\eta_t \eta_t' \mid t \in \mathcal{R}_2].$$

**Identification** This form of heteroskedasticity, together with the assumption that the  $n \times n$  structural impact matrix  $A$  remains unchanged across regimes, implies the following moment conditions:

$$\Omega_1 = A\Sigma_1 A' \tag{G.1a}$$

$$\Omega_2 = A\Sigma_2 A' \tag{G.1b}$$

where  $\Omega_i = \mathbb{E}[x_t x_t' \mid t \in \mathcal{R}_i]$  is the covariance matrix of the observable variables  $x_t$  in regime  $i$ . Thus, given estimates of  $\Omega_1$  and  $\Omega_2$  we can use the above moment conditions to identify  $A$ ,  $\Sigma_1$ , and  $\Sigma_2$ .

The each of the symmetric matrices  $\Omega_1$  and  $\Omega_2$  each provide  $n(n + 1)/2$  empirical moments. With the normalization that  $\Sigma_1 = I_n$  (which is without loss of generality), there are  $n^2$  parameters to identify in  $A$  and  $n$  in  $\Sigma_2$ . Putting these together, there are  $n(n + 1)$  empirical moments and  $n(n + 1)$  parameters to estimate, so the system is exactly identified.

**Estimation** In practice, estimation of the moment conditions in equation (1.12) is complicated by the fact that ordering and sign of the columns (and elements of  $\Sigma_2$ ) are not identified. This does not present a challenge for point estimates, but it does present a challenge for inference. Thus, while a typical GMM estimation would rely on a numerical optimizer for estimation, the optimizer has no way of keeping “the same” shocks in the same order.

A numerical optimizer can, however, be circumvented (which is also nice for stability reasons). Notice that (again, maintaining the assumption that  $\Sigma_1 = I_n$ )

$$\Omega_2\Omega_1^{-1} = A\Sigma_2A'(AA')^{-1} = A\Sigma_2A^{-1}.$$

The final expression resembles an eigendecomposition of  $\Omega_2\Omega_1^{-1}$ . In fact, any eigendecomposition of  $\Omega_2\Omega_1^{-1}$  can be used as a starting point for estimation—this is convenient since fast and stable algorithms exist for performing these decompositions.

Specifically, let  $Q$  and  $\Lambda$  form an eigendecomposition of  $\Omega_2\Omega_1^{-1}$ :

$$\Omega_2\Omega_1^{-1} = Q\Lambda Q^{-1}. \tag{G.2}$$

Eigendecompositions are not in general unique; while the eigenvalues (the diagonal elements of  $\Lambda$ ) are unique (up to ordering—an issue that is of no consequence for point estimates but will return when conducting inference), each eigenvector can be multiplied by a different scalar and the decomposition will be preserved.<sup>1</sup> However, information from the first regime can be used to

---

<sup>1</sup>Let  $D$  be diagonal and  $\tilde{Q} \equiv QD$ . Diagonal matrices commute, so  $\tilde{Q}$  and  $\Lambda$  form another eigendecomposition:

$$\tilde{Q}\Lambda\tilde{Q}^{-1} = QD\Lambda D^{-1}Q^{-1} = Q\Lambda DD^{-1}Q^{-1} = Q\Lambda Q^{-1} = \Omega_2\Omega_1^{-1}.$$

pin down the unique scaling of the eigenvectors such that the moment conditions are satisfied. Set  $A = QS$  where

$$S \equiv \sqrt{\text{diag} [Q^{-1}\Omega_1Q'^{-1}]},$$

where the square root operates on each element of the embedded diagonal matrix. Then

$$AA' = QSS'Q' \tag{G.3a}$$

$$= Q\text{diag} [Q^{-1}\Omega_1Q'^{-1}] Q' \tag{G.3b}$$

$$= \Omega_1 \tag{G.3c}$$

see proposition 2.

Letting  $\Sigma_2 = \Lambda$ , we have

$$A\Sigma_2A' = Q\Lambda S'Q'$$

$$= Q\Lambda SS'Q' \tag{diagonal matrices commute}$$

$$= Q\Lambda\text{diag} [Q^{-1}\Omega_1Q'^{-1}] Q'$$

$$= Q\Lambda\text{diag} [\Lambda^{-1}Q^{-1}\Omega_2Q'^{-1}] Q' \tag{using (G.2)}$$

$$= Q\Lambda\Lambda^{-1}\text{diag} [Q^{-1}\Omega_2Q'^{-1}] Q' \tag{diagonal matrices commute}$$

$$= \Omega_2 \tag{see proposition 2.}$$

Thus, the moment conditions in equations (1.12) are satisfied by  $A = QS$ ,  $\Sigma_1 = I$ , and  $\Sigma_2 = \Lambda$ .

As a brief aside, the fact that the heteroskedasticity-based identification (HBI) estimate of  $\mathbf{M}$  consists of the eigenvectors of the ratio of the empirical covariance matrices  $\Omega_2\Omega_1^{-1}$  reveals a connection to principal components analysis (PCA). To the best of my knowledge this relationship has not been described previously. The PCA estimate of  $\mathbf{M}$  (factor loadings) consists of the eigenvectors of the covariance matrix of the observable variables—this covariance matrix is given by the full-sample  $\Omega$ . Intuitively, this implies that PCA factors are designed to explain the largest amount

of variation in the data with the fewest number of factors. By analogy, this means that the shocks estimated by HBI are those that can best explain the *relative* variance of the shocks in the two regimes, since they are based on the eigendecomposition of the *relative* covariance matrices.

**Inference** To conduct inference on the parameters estimates above (and functions of those estimates) I rely on a bootstrapping procedure. I begin by drawing observations from the two regimes with replacement, stratifying by the size of the regimes. Thus, I rely on a completely standard bootstrap, save for one computational difficulty. The estimation procedure above did not identify the ordering or sign of the columns of  $A$ .<sup>2</sup>

In order to increase the likelihood that the shocks from each bootstrap sample are in the same order and sign as the shocks that form my point estimates, I rely on an “aligning” procedure similar to that laid out in Clarkson (1979). Let  $(A, \Sigma_2)$  be the point estimates of the system and  $(A^b, \Sigma_2^b)$  be the estimates from a particular bootstrap sample  $b$ . I search over the set of all column permutation matrices  $\mathcal{P}$  and the set  $\mathcal{S}$  of  $n$ -dimensional diagonal matrices with elements in  $\{-1, 1\}$  and define

$$\begin{aligned} (A^{b*}, \Sigma_2^{b*}) &= \underset{P \in \mathcal{P}, S \in \mathcal{S}}{\operatorname{argmin}} (1 - \lambda) \left\| \tilde{A}^b - A \right\|_F + \lambda \left\| \tilde{\Sigma}_2^b - \Sigma_2 \right\|_F \\ \text{s.t. } \tilde{A}^b &= A^b P S \\ \tilde{\Sigma}_2^b &= \Sigma_2^b P \end{aligned}$$

for some  $\lambda \in [0, 1]$ . The elements  $\{A^{b*}, \Sigma_2^{b*}\}_b$  form my bootstrap distribution. In words, I rearrange the order and sign of the columns of  $A^b$  (and the corresponding order of the elements of  $\Sigma_2^b$ ) such that the distance between the rearranged matrices are closest (under the Frobenius norm) to the point estimates. In practice the  $\lambda$  matters little, as long as it is interior. I calculate confidence intervals using the percentiles of the bootstrap distribution.

I conclude the section by proving a matrix equality used earlier in this appendix.

---

<sup>2</sup>The eigendecomposition did not identify the order, and any version of  $A$  with flipped column signs would have satisfied the moment conditions since these cancel out when taking  $AA'$ .

**Proposition 2** *In the context of equation (G.3), the following equality holds:  $Q \text{diag} [Q^{-1} \Omega_1 Q'^{-1}] Q' = \Omega_1$  as long as the elements of  $\Lambda$  are unique (i.e.  $\Lambda_{ii} \neq \Lambda_{jj} \forall i \neq j$ ).<sup>3</sup>*

*Proof.*

We first need access to the following relation for a diagonal matrix  $D$  with distinct diagonal  $d_i$  and matrix with unknown properties  $X$ :

$$DXD^{-1} = X \implies X \text{ is a diagonal matrix.} \quad (\text{G.4})$$

To see this, suppose that for  $i \neq j$ ,  $x_{ij} \neq 0$ . Note that the  $ij^{\text{th}}$  element of  $DXD^{-1}$  is given by  $d_i d_j^{-1} x_{ij}$ . By assumption,  $d_i d_j^{-1} x_{ij} = x_{ij}$ . Since  $x_{ij} \neq 0$ , we can divide both sides by  $x_{ij}$  to get  $d_i = d_j$  which contradicts the assumption that the diagonal elements of  $D$  are unique.

This result allows us to see that  $Q^{-1} \Omega_1 Q'^{-1}$  is diagonal. Performing two rearrangements of equation (G.2), we have

$$\begin{aligned} Q'^{-1} &= (Q^{-1})' = (\Lambda^{-1} Q^{-1} \Omega_2 \Omega_1^{-1})' = \Omega_1^{-1} \Omega_2 Q'^{-1} \Lambda^{-1} \\ Q^{-1} &= (\Omega_2 \Omega_1^{-1} Q \Lambda^{-1})^{-1} = \Lambda Q^{-1} \Omega_1 \Omega_2^{-1}. \end{aligned}$$

With these we can see that

$$\begin{aligned} Q^{-1} \Omega_1 Q'^{-1} &= \Lambda Q^{-1} \Omega_1 \Omega_2^{-1} \Omega_1 \Omega_1^{-1} \Omega_2 Q'^{-1} \Lambda^{-1} \\ &= \Lambda Q^{-1} \Omega_1 Q'^{-1} \Lambda^{-1} \end{aligned}$$

which, by equation (G.4), implies that  $Q^{-1} \Omega_1 Q'^{-1}$  is diagonal.

---

<sup>3</sup>Note that this condition on  $\Lambda$  is implied by the heteroskeasticity-based assumptions that no two shocks change in the same proportion.

Next we establish notation. Let  $E_k$  be an  $n \times n$  matrix whose elements  $e_{ij}$  are given by

$$e_{ij} = \begin{cases} 1 & i = j = k \\ 0 & \text{otherwise.} \end{cases}$$

For intuition,  $E_k X$  zeros out all but the  $i^{\text{th}}$  row of  $X$ , and  $X E_k$  zeros out all by the  $i^{\text{th}}$  column of  $X$ .

Note that for a diagonal matrix  $D$ , when  $i \neq j$ ,

$$E_i D E_j = E_j D E_i = 0_n \quad (\text{G.5})$$

an  $n \times n$  matrix of zeros. Note further that

$$\text{diag}(X) = \sum_{i=1}^n E_i X E_i.$$

With this notation and the above-established result, we can proceed to the following derivation.

$$\begin{aligned} Q \text{diag} \left[ Q^{-1} \Omega_1 Q'^{-1} \right] Q' &= Q \left[ \sum_i E_i Q^{-1} \Omega_1 Q'^{-1} E_i \right] Q' \\ &= \sum_i Q E_i Q^{-1} \Omega_1 Q'^{-1} E_i Q' \\ &= \left[ \sum_i Q E_i \right] Q^{-1} \Omega_1 Q'^{-1} \left[ \sum_i E_i Q' \right] \quad \text{by (G.5), since } Q^{-1} \Omega_1 Q'^{-1} \text{ is diagonal} \\ &= Q Q^{-1} \Omega_1 Q'^{-1} Q' \quad \text{definition of } E_i \\ &= \Omega_1, \end{aligned}$$

which is what we wanted to show.


*Q.E.D.*

## Appendix H: Local Projections: Robustness


This section contains several variations on the choices made to estimate equation (1.17).

### 1. Adding months with no FOMC meetings (setting shocks to 0)

#### Responses to Contractionary Monetary Path Shock ( $\beta_k^M$ )




#### Responses to Expansionary Information Shock ( $\beta_k^I$ )




**2. Sample period** Stopping the estimation in 2015 ensures the same number of observations for each regression horizon, and drops COVID observations). The “Keep GR” response does not drop the 07/08–07/09 dates.

#### Responses to Contractionary Monetary Path Shock ( $\beta_k^M$ )




Responses to Expansionary Information Shock ( $\beta_k^I$ )




3. More (24) and fewer (0) lags as controls; adding a trend

Responses to Contractionary Monetary Path Shock ( $\beta_k^M$ )




Responses to Expansionary Information Shock ( $\beta_k^I$ )




#### 4. Using the ten-word word list

##### Responses to Contractionary Monetary Path Shock ( $\beta_k^M$ )






##### Responses to Expansionary Information Shock ( $\beta_k^I$ )






#### 5. Tri-variate high-frequency system in Eurodollar futures, output forecast revisions, and the surprise in the Federal Funds rate, from section 1.5.1




##### Responses to Contractionary Monetary Path Shock ( $\beta_k^M$ )



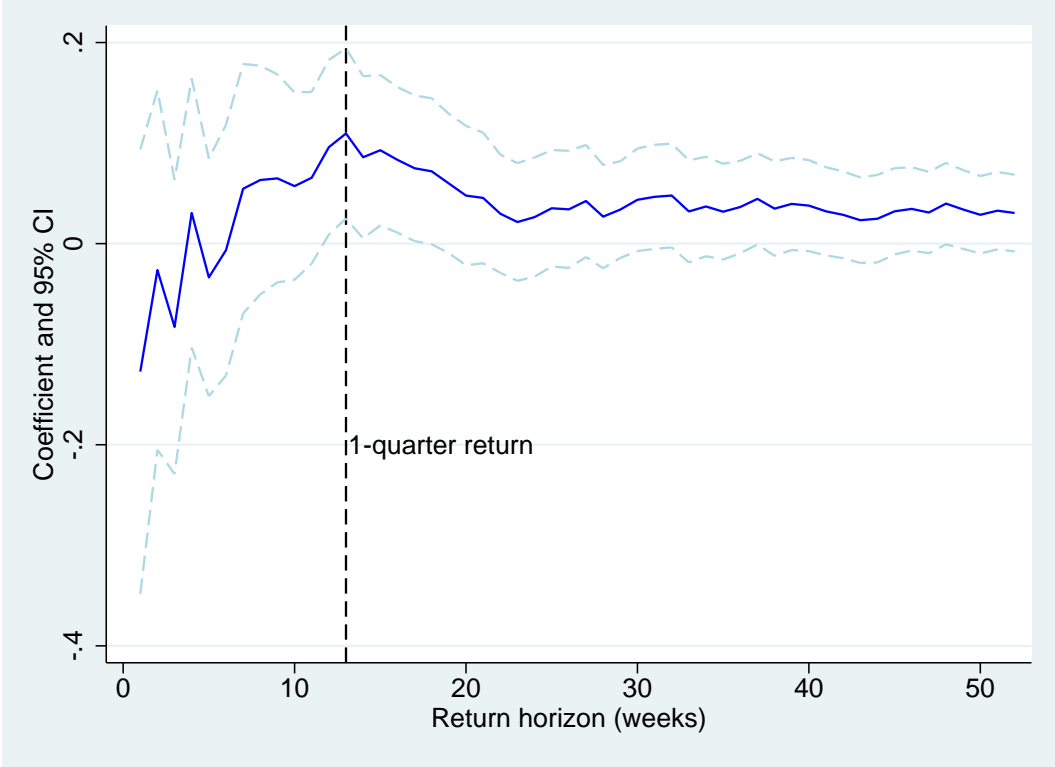
Responses to Expansionary Information Shock ( $\beta_k^I$ )



Responses to Contractionary Target Shock



**Appendix I: Predictability of Interest-Rate Surprises: A Comment on Bauer and Swanson (2020)**



This figure presents estimates of the following regression

$$\hat{i}_t^{\text{NS}} = \alpha_h + \beta_h(p_{t-1} - p_{t-7h}) + e_t$$

where  $\hat{i}_t^{\text{NS}}$  is the high-frequency interest-rate surprise of Nakamura and Steinsson (2018) (from their replication materials), and  $p_t$  is the log of the S&P 500 index on day  $t$ 's market close. The sample consists of all days ( $t$ ) with regularly-scheduled FOMC meetings between 1995 and 2015 (excluding the July 2008–July 2009 period) and after the first week of the month (i.e., the observations used when testing for the presence of information effects with Blue Chip data). The

right-hand-side variable,  $p_{t-1} - p_{t-7h}$ , is thus the  $h$ -week return in the S&P 500 ending the day before each FOMC meeting. I estimate the regression for each  $h$ , which is shown on the  $x$ -axis. Bauer and Swanson (2020) present results using the 13-week return, which I highlight with a dashed line. This return horizon has the largest and most statistically significant coefficient.

## Appendix J: Assessing the Empirical Assumptions

In this section, I present additional evidence pertaining to two central assumptions made in the specification of the empirical model of section 1.3.3: the independence of the *perceived* structural shocks, and the stability of the  $m_{ij}$  coefficients. A useful extension of the model will aid, statistically, in backing up these assumptions: the estimation of an over-identified system wherein I modify the form of the assumed heteroskedasticity to

$$\sigma_{1,t} = \begin{cases} 1 & t \in \mathcal{R}_1 \\ \sigma_1 & t \in \mathcal{R}_2 \\ \tilde{\sigma}_1 & t \in \mathcal{R}_3 \end{cases} \quad \sigma_{2,t} = \begin{cases} 1 & t \in \mathcal{R}_1 \\ \sigma_2 & t \in \mathcal{R}_2 \\ \tilde{\sigma}_2 & t \in \mathcal{R}_3, \end{cases} \quad (\text{J.1})$$

where  $\mathcal{R}_3$  covers the rest of the sample: from January 2007 through December 2019. In section J.1, I assess the independence assumption using statistical evidence from the over-identified model, and narrative evidence from newspaper articles. In section J.2, I address the stability assumption in three ways: first, statistically, using the over-identified mode; second, empirically, using an estimated DSGE model in which the assumption does not hold; third, I discuss theoretical conditions that invalidate the assumption.

### J.1 Uncorrelated Forecast Revisions

**Narrative Evidence** [TBD]

**Statistical Evidence** The additional regime of equation (J.1) brings in three more moments in the data but only two additional parameters: the variances of the structural shocks in the third regime,  $\tilde{\sigma}_1$  and  $\tilde{\sigma}_2$ . I take advantage of the remaining moment to allow for correlation between the

Table J.1: Allowing for Correlation in Perceived Shocks

	Correlation w/ Baseline Shocks			
	$\rho$	(s.e.)	Monetary	Information
Common: Regimes 1 and 2	-0.121	(0.306)	0.999	0.996
Regime 1 Only	0.317	(0.279)	0.797	0.807
Regime 2 Only	-0.095	(0.242)	0.993	0.976

structural shocks with the introduction of the correlation coefficient:

$$\rho_t \equiv \frac{\mathbb{E}[\xi_{1,t}\xi_{2,t}]}{\sigma_{1,t}\sigma_{2,t}} = \begin{cases} \rho_{(1)} & t \in \mathcal{R}_1 \\ \rho_{(2)} & t \in \mathcal{R}_2 \\ 0 & t \in \mathcal{R}_3 \end{cases}.$$

I consider three assumptions for  $\rho_t$ .<sup>1</sup> First, I allow the correlation to be non-zero, but constant, in the first two regimes (mathematically,  $\rho_{(1)} = \rho_{(2)}$ ). Second, I allow non-zero correlation in the first regime only ( $\rho_{(2)} = 0$ ), arguably the regime most-susceptible to correlated forecast revisions given the *relatively* low volume of communications. Third, for completeness, I allow for non-zero correlation in the second regime only ( $\rho_{(1)} = 0$ ).

Table J.1 shows the results. The first column shows the estimates of  $\rho_t$ , with its standard error alongside. The third and fourth columns show the correlation of the perceived shocks estimated under each alternative specification with the baseline shock series. In all cases  $\rho_t$  is imprecisely estimated. Its magnitude is fairly small when the second regime is included, but larger in the first regime, consistent with the notion that Fed watchers likely found it harder to disentangle signals about the fundamental shocks when the Fed did not explicitly communicate about future monetary policy. Despite this correlation, the estimated perceived shocks are quite similar to the baseline estimates, with correlation coefficients of 0.8.

<sup>1</sup>In all cases, I enforce that the perceived shocks are uncorrelated in the third regime ( $\rho_3 = 0$ ). This is for practical purposes (with non-zero correlation in all regimes, the system is rank deficient), but also empirically relevant given the vast amount of communications that the Fed has released with its announcements since 2008.

## J.2 Constant Mapping between Shocks and Observables

### J.2.1 Statistical Validity

Statistically, the over-identified model provides a natural framework in which to test for parameter stability using tests of overidentifying restrictions. I carry out two tests that, taken together, support the assumption of parameter stability. First, the Hansen (1982) J-test for overidentifying restrictions fails to reject those restrictions ( $p = 0.68$ ). Second, the weak-instrument robust K-test of Kleibergen (2005) fails to reject the joint null hypothesis that the  $m_{ij}$  coefficients change from the exactly- to over-identified system ( $p = 0.89$ ).<sup>2</sup>

### J.2.2 Empirical Validity: Melosi (2017)

In this appendix, I explore whether the assumptions I made to identify structural shocks in my empirical application allow me to identify the relevant structural shocks in a theoretical context in which those assumptions are known not to hold. To that end, I study an extension of the model of Melosi (2017).

**Economy** The log-linear model equations are

$$y_t = \mathbb{E}_t[y_{t+1}] - (i_t - \mathbb{E}_t[\pi_{t+1}]) + (g_t - \mathbb{E}_t[g_{t+1}]) \quad (\text{IS})$$

$$i_t = \phi_\pi (\pi_t + \xi_{\pi,t}) + \phi_x (x_t - a_t + \xi_{x,t}) + \xi_{m,t} \quad (\text{TR})$$

$$\pi_t = (1 - \theta)(1 - \beta\theta) \sum_{k=1}^{\infty} (1 - \theta)^{k-1} \text{mc}_{t|t}^{(k)} + \beta\theta \sum_{k=1}^{\infty} (1 - \theta)^{k-1} \pi_{t+1|t}^{(k)}. \quad (\text{ICKPC})$$

where  $\text{mc}_{t|t}^{(k)} \equiv y_{t|t}^{(k)} - a_{t|t}^{(k-1)}$ . The model is evidently similar to the standard three-equation New-Keynesian model, describing the evolution of GDP ( $y_t$ ), interest rates ( $i_t$ ), and inflation ( $\pi_t$ ) according to, in order, an IS equation (IS), Taylor Rule (TR), and Phillips curve (ICKPC). The model's

---

<sup>2</sup>Here, I fix the relative variances  $\sigma_1$  and  $\sigma_2$  at their exactly-identified values and test for whether the  $m_{ij}$  coefficients jointly change.

fundamental shocks are the household’s discount factor shock (the “demand” shock,  $g_t$ ), an aggregate productivity shock ( $a_t$ ), an exogenous monetary shock ( $\xi_m$ ), and two shocks reflecting the monetary authority’s imperfect measurement of inflation and the output gap ( $\xi_{\pi,t}$  and  $\xi_{x,t}$ , respectively). All shocks evolve according to mutually uncorrelated first-order autoregressive processes. The IS equation is standard, determined by perfectly-informed households’ optimal consumption-saving decision and the equilibrium condition that firms must meet household demand in each period, having set prices in the beginning of each period. The Taylor rule is mostly standard, save for the Fed’s imperfect reading of the state of the economy. The imperfect common knowledge Phillips curve is determined by firms’ (indexed by  $j$ ) optimal pricing behavior, based on their (imperfect) information sets ( $\mathcal{I}_{j,t}$ , to be described shortly). Because firms are monopolistically competitive, each firm’s optimal price is a function of aggregate demand *and* the average price of their competitors. As introduced and described by Woodford (2003), not having access to the their competitors’ information sets, firms are left not only to form expectations of the current state of the economy, but also about their competitors’ expectations which in turn depend on their competitors’ expectations of their competitors’ expectations, and so on. Thus the notation for these “higher-order expectations,” used above, for a generic variable  $z_t$ :

$$z_{t|t}^{(0)} = z_t \quad z_{t|t}^{(1)} = \int \mathbb{E}_{j,t} [z_t^{(0)}] d_j = \int \mathbb{E}_{j,t} [z_t] d_j \quad \dots \quad z_{t|t}^{(k)} = \int \mathbb{E}_{j,t} [z_t^{(k-1)}] d_j.$$

Of particular relevance for my empirical application is  $z_{t|t}^{(1)}$ —average expectations about  $z_t$ .

**Information** The information structure for a firm  $j$  is thus central for determining the behavior of aggregate variables. If, as traditionally assumed, firms were perfectly informed, the economy’s Phillips curve would collapse to the familiar  $\pi_t = \kappa mc_t + \beta \mathbb{E}_t[\pi_{t+1}]$ . Firms are not perfectly informed, however, and use the Kalman filter to update expectations based on the information available to them.<sup>3</sup> I assume that each discrete period, indexed by integers  $t$ , is split into the

---

<sup>3</sup>This information structure is not necessarily the optimal one that firms would chose if they had a constraint on the extent to which they can process information. If we assume, however, that firms “in the real world” collect information in this way, then it may be reasonable to assume that Melosi’s estimates capture the optimal values of

following segments:

- at time  $t$ , all shocks are realized, and the Fed sets interest rates;
- at time  $t_F$ , the Fed makes its policy announcement, endowing firms with independent signals of the monetary ( $\xi_m$ ) and demand ( $g_t$ ) shocks, each buffeted by Gaussian noise, and firms update their expectations;
- at time  $t_P$ , firms observe their private productivity (a Gaussian deviate of aggregate productivity) and a private signal about demand, update their expectations, and set prices; and
- at time  $t_E$ , the representative household becomes perfectly informed and makes its consumption/savings decision, and firms produce to meet household demand

where  $t < t_F < t_P < t_E < t + 1$ . This setup follows that of Melosi, with two notable modifications. First, Melosi assumes that firms observe the Fed’s announcement and their private signals simultaneously. Splitting these phases up allows me to measure how expectations are revised following the Fed’s policy announcements. Second, more fundamentally, I make the Fed’s announcements more-detailed than Melosi, who construes of those announcements as simply the revelation of  $i_t$ . At least within the context of my 1999–2019 empirical sample, that is an outdated representation. The specific shocks that I assume are signaled at  $t_F$  are consistent with my empirical findings—markets appear to learn primarily about demand and monetary shocks, and are able to distinguish the information provided about both shocks.<sup>4</sup>

**Calibration** My baseline calibration closely follows Melosi’s estimates. While those estimated values speak to the *totality* of the information firms have, they cannot speak to the *source* of the information. By splitting up firm’s information streams into the Fed announcement and private


---

parameters chosen by such rationally inattentive firms.


<sup>4</sup>If firms only observed  $i_t$ , they would have to infer which fundamental shock (demand, technology, or monetary policy) drove the interest rate decision, based on other private information they observe and their priors. Forecast revisions of these (mutually uncorrelated) structural shocks would therefore be correlated. My empirical results suggest that these forecast revisions are essentially uncorrelated (see figure 6).

Figure J.1: Heteroskedasticity in the Model

Panel A: Baseline Model



Panel B: Common Knowledge Model



revelation, I thus have to take a stance on what information is learned when. I thus assume, as estimated by Melosi, that firms learn very little about demand from their private signal (signal-to-noise ratio of  $\approx 10$ ) whereas Fed announcements reveal demand with twice the precision of those private signals. My baseline calibration for the signal-to-noise ratio of the monetary shock signal is about ten. I chose these values to broadly match my data on high-frequency forecast revisions.<sup>5</sup>

**Experiment** Ultimately I am interested in measuring the extent to which the shocks identified using the heteroskedasticity-based identification (HBI) assumptions resemble their theoretical counterparts. To assess this, I simulate the model under two regimes. For 5000 periods, I simulate the model using the baseline calibration described above. I then reduce the variance of the noise with

---

<sup>5</sup>A more-satisfactory approach would be to measure, in Melosi's model, the mutual information that firms' information sets provide about each structural shock. I could then keep this level of mutual information fixed and chose how to partition it between  $t_F$  and  $t_P$ .

which the Fed communicates about the monetary shock by a third, and simulate the model for 5000 more periods. This is meant to capture the “introduction of forward guidance” that I study empirically.<sup>6</sup> In each simulated series, I calculate forecast revisions made about the model’s endogenous variables and exogenous shocks around the Fed’s announcement.<sup>7</sup> Denote the changes in expectations about output and interest rates by  $\hat{y}_t$  and  $\hat{i}_t$ . Note that firms only update expectations about the demand and monetary shocks following Fed announcements, since these are the only shocks that the Fed discusses.

Next, I append the two sets of simulated data and estimate two shocks using the HBI assumptions, applied to the vector of observable variables  $[\hat{y}_t, \hat{i}_t]$ . These are plotted in panel A of figure J.1. Recall that the model’s endogenous variables—including output and interest rates—will be a function of not only exogenous shocks, but increasingly higher-order expectations of those exogenous shocks. Thus, forecast revisions of  $y_t$  and  $i_t$ —and the resulting HBI shocks—will be linear combinations of forecast revisions of not only demand and monetary shocks, but also higher-order expectations thereof. The question is: *how much* of the variation in each HBI shock comes from forecast revisions about each structural shock? Because I can directly measure forecast revisions about  $\{\xi_{m,t}, g_t, \xi_{m,t}^{(1)}, g_t^{(1)}, \xi_{m,t}^{(2)}, g_t^{(2)}, \dots\}$ , I project each HBI shock onto all of these forecast revisions, and can thus decompose the variance of each shock into (1) the amount explained by forecast revisions about the monetary shock (and higher-order expectations thereof), (2) the amount explained by forecast revisions about the demand shock (and higher-order expectations thereof), and (3) any covariance between (1) and (2).<sup>8</sup>

Figure J.2 shows the variance decomposition of the “reduced-form perceived monetary shock”


---

<sup>6</sup>Since the monetary shock is highly persistent in this model, this is a reasonable proxy for explicit forward guidance, since it signals future deviations from the systematic component of the Taylor rule.

<sup>7</sup>Formally, for a variable  $\eta_t$ , I calculate  $\eta_{t|tF}^{(1)} - \eta_{t|(t-1)P}^{(1)}$ —the change in the average forecast—consistent with my data.

<sup>8</sup>Formally, let  $\hat{s}_t$  be an HBI shock. I estimate the OLS regression  $\hat{s}_t = \underbrace{\sum_{k=0}^{10} \phi_{m,k} \hat{\xi}_{m,t}^{(k)}}_{\text{monetary factors}} + \underbrace{\sum_{k=0}^{10} \phi_{g,k} \hat{g}_t^{(k)}}_{\text{demand factors}}$ . (Note that this projection has no error term.) I can then use the estimated  $\phi$  coefficients to decompose the variance of  $\hat{s}_t$  into the variance arising from monetary factors, the variance arising from demand factors, and the covariance between monetary factors and demand factors.

Figure J.2: Composition of the Heteroskedasticity Estimator



identified by heteroskedasticity (the shock with the larger variance in the second regime) and the “reduced-form information shock.” In the baseline model, it is apparent that these HBI shocks uncover their structural counterparts remarkably well. The information shock is primarily made up of forecast revisions about the structural demand shock, while the perceived monetary shock primarily measures the structural monetary shock. That the covariance term is negligible is expected given the fact that firms observe independent signals about each structural shock.

**Discussion: The Common Knowledge Benchmark** To understand why the shocks identified with the HBI assumptions do not perfectly measure the model’s structural shocks, it is useful to consider a variant of the baseline in which firms have common-knowledge, only forming expectations based on public signals.<sup>9</sup> There, the HBI and structural shocks perfectly coincide. This coincidence is verified in figure J.2, where the HBI shocks were estimated using the data in panel B of figure J.1.<sup>10</sup> To see why the shocks are correctly identified in this model, it is useful to

<sup>9</sup>The common-knowledge model therefore resembles that of Jia (2020). Currently I have implemented the common-knowledge model by replacing (ICKPC) with  $\pi_t = \kappa mc_{t|t} + \beta \pi_{t+1|t}$  as in Melosi’s perfect information model. Alternatively, I could follow Jia exactly and replace (ICKPC) with  $\pi_t = \beta \theta \pi_{t+1|t} + (1 - \theta) \pi_{t|t} + \kappa \theta \hat{y}_t$ . For reasons I discuss below, this will not change my conclusions—HBI will still identify forecast revisions of the corresponding structural shocks.

<sup>10</sup>I make a few modifications to the baseline parameterization so that the simulated data continue to resemble the data from the baseline model. I shut off  $\{a_t, \xi_{\pi,t}, \xi_{x,t}\}$ , though this is likely not crucial and something to revisit. More pertinently, I reduce the autocorrelation of the monetary shock from 0.94 to 0.3—similar to Melosi’s estimate of this parameter in his perfect information model. Without this, positive monetary shocks lead to decreases in  $i_t$  (an expected result—see Galí (2015)), so monetary and demand shocks “look” the same in terms of the sign of their impact responses on  $i_t$  and  $y_t$ . Finally, I reduce the signal-to-noise ratio of the monetary shock from 1.25 to 0.25.

consider how the model is solved. The solution of the model is assumed to take the form

$$x_t = M\eta_{t|t} \tag{J.2}$$

where  $x_t$  are the model's endogenous variables ( $y_t, i_t$ , and  $\pi_t$ ), and  $\eta_{t|t}$  are agents' perceptions of the model's structural shocks, and high-order expectations of those shocks.<sup>11</sup> Those expectations are assumed to follow an VAR(1) process

$$\eta_{t|t} = L\eta_{t-1|t-1} + N\xi_t \tag{J.3}$$

where  $\xi_t$  are innovations to the economy's structural shocks. The matrices  $M, L$ , and  $N$  form the model's solution, and are found by iterating back and forth between equations (J.2) and (J.3) until convergence. If firms do not observe endogenous variables, the solution method is even simpler: solve for  $L$  and  $N$  by solving firms' Kalman filtering problem, then solve for  $M$  a solver for linear rational expectations models (see Appendix C).

Writing equation (J.2) in expectations-revisions space allows us to see that forecast revisions about observable variables are linear combinations about structural shocks:

$$\hat{x}_t = M\hat{\eta}_{t|t}.$$

The question is therefore: does changing the noise with which signals are communicated affect  $M$ ? In the common-knowledge model, the answer is no. The change in signal clarity affects firms' abilities to infer signals, but this does not change anything fundamental about their response to a *perceived shock of the same magnitude*. Instead, in the baseline (imperfect common knowledge) model, expectation formation plays a fundamental role, altering the relationship between firm  $j$ 's price and (perceived) aggregate demand. In that model, the presence of idiosyncratic signals causes firms to remain confused as to the source of those signals, unclear as to whether they reflect ag-

---

<sup>11</sup>Thus, in principle  $\eta_{t|t}$  is infinite-dimensional. Following Melosi, I cap higher-order expectations are at 10.

gregate or idiosyncratic shocks. This results in firms placing different weights on expectations of different higher-orders when the precision of their signals changes, altering the “slope” of the Phillips curve. It is this sense in which the change is “fundamental” in the baseline model. The next section explores these issues in more detail.

### J.2.3 Dispersed vs. Common Knowledge

The discussion in the text, and in the previous appendix, noted that models with dispersed information do not feature a constant mapping between forecast revisions of observable variables, and forecast revisions of structural shocks. In this appendix I analyze a simple model to explain this point. Note that the model in the previous appendix featured dispersed information. However, I showed that the shocks identified by heteroskedasticity uncovered their structural counterparts remarkably well.

I consider a model of the evolution of an asset price  $p_t$ , whose determination depends on the average expected future price of the asset and an exogenous fundamental:

$$p_t = \beta \int \mathbb{E}[p_{t+1} \mid \Omega_{j,t}] dj + \theta_t \quad \theta_t = \rho \theta_{t-1} + \zeta_t \quad \zeta_t \sim N(0, \sigma_\zeta^2) \quad (\text{J.4})$$

where  $\Omega_{j,t}$  is the information set of agent  $j$  for  $j \in (0, 1)$ , and  $\beta, |\rho| \in [0, 1)$ .<sup>12</sup> All agents receive mutually independent signals  $s_{t,j}$  regarding the fundamental, which they never forget:

$$s_{t,j} = \theta_t + v_{t,j} \quad v_{t,j} \sim \mathcal{N}(0, \sigma_v^2) \quad \Omega_t = \Omega_{t-1} \cup \{v_{t,j}\}.$$

Introducing the notation  $x_{t+j|t}^{(k)} = \int x_{t+j|t}^{(k-1)} dj$  with  $x_{t+j|t}^{(0)} = x_t$ , the model can be expressed by recursive substitution as

$$p_t = \sum_{k=0}^{\infty} \beta^k \theta_{t+k|t}^{(k)}. \quad (\text{J.5})$$

---

<sup>12</sup>This dynamic asset pricing problem shares features with to models studied by Townsend (1983), Morris and Shin (2006), and Woodford (2003). Nimark (2017) is a fantastic reference for understanding models of this form. The author proposes a solution method for generalized models of the form I present in equation (J.4), a special case of which is the model of Melosi (2017).

The question of relevance for my empirical results is whether a change in  $\nu_{t,j}$  (or  $\sigma_\zeta^2$ ) changes the mapping between forecast revisions regarding  $p_t$ ,  $p_{t|t}^{(1)} - p_{t|t-1}^{(1)}$ , forecast revisions regarding  $\zeta_t$ ,  $\theta_{t|t}^{(1)} - \theta_{t|t-1}^{(1)}$ . The fairly complex form of (J.5) (specifically, the  $\theta$  terms therein) suggests that this need not be the case. However, in the common knowledge case, with  $\nu_{t,j} = \nu_{t,i}$ ,  $\forall(i, j)$ , the mapping is constant. To see this, notice that in this case the model reduces to

$$p_t = \sum_{k=0}^{\infty} \beta^k \theta_{t+k|t}^{(1)} = \sum_{k=0}^{\infty} \beta^k \rho^k \theta_{t|t}^{(1)} = \frac{1}{1 - \beta\rho} \theta_{t|t}^{(1)}$$

so that forecast revisions are

$$p_{t|t}^{(1)} - p_{t|t-1}^{(1)} = \frac{1}{1 - \beta\rho} \left[ \theta_{t|t}^{(1)} - \theta_{t|t-1}^{(1)} \right],$$

where the mapping  $\frac{1}{1 - \beta\rho}$  evidently does not depend on the variance of the fundamental or noise.<sup>13</sup>


---

<sup>13</sup>This is not to say that the relationship between the fundamental itself and the price level remains unchanged. Setting  $\rho = 0$ , the model's solution is  $p_t = \left( \frac{1}{1 - \beta\rho} \right) \left( \frac{\sigma_\zeta^2}{\sigma_\zeta^2 + \sigma_\nu^2} \right) (\theta_t + \nu_t)$ . The conclusion without  $\rho = 0$  reveals a solution whereby  $p_t$  is a function of all realized fundamental and noise shocks, with a mapping onto prices that also depends on the variances of those shocks.


## Appendix K: Weighted Statistics

This appendix represents the average ad-valorem equivalent MFN tariff on low- and high-value varieties, weighted by each variety's (left) or good's (right) import share. Goods are split by median unit-value, as in the baseline. The results are similar to the unweighted average differences, presented as the baseline estimates in the paper.

Variety Weigting




Good Weighting



## Appendix L: Regressivity in the European Tariff Schedule

This appendix shows the analogue of our main U.S. figures using the 2017 EU tariff schedules. The panel on the left shows the fraction of regressive goods within the full, and various subsets of the EU tariff schedules. We define a “good” as an HS5/unit of quantity pair, and restrict our analysis to only ad valorem rates of duty (the more conservative approach to checking for regressivity). The panel on the right shows the tariff differential among varieties of regressive goods.



## Appendix M: Details on Digitized Tariff Data

This appendix describes the data sources used to form our new dataset of legislated tariff rates. Table M.1 shows the list of documents that we digitized. The first document, *Foreign Commerce and Navigation of the U.S.*, provides tariff rates immediately before and after the Smoot-Hawley Act. All legislated tariff changes between 1931 and 1946 were included in the first edition of the Tariff Commission’s *United States Import Duties* (USID). Between 1946 and 1958 we gathered the annual USID after each round of GATT negotiations (except the first Geneva round). The year 1963 witnessed a major overhaul of the tariff schedules, which culminated in the creation of the *Tariff Schedules of the United States* (TSUS). We digitized the TSUS in the year following each GATT round from 1963 through 1987 (Dillon, Kennedy, and Tokyo). The year 1989 saw another major overhaul of the import classification, with the establishment of the Harmonized Tariff Schedules (HTS). Starting at this point, digitized tariff schedules are readily available.

Table M.1: Data Sources

Year	Publication Name	Event
1930	<i>Foreign Commerce and Navigation of the U.S.</i>	Pre- and Post-Smoot Hawley
1931–1946	<i>United States Import Duties</i>	Pre-Geneva I
1948	<i>United States Import Duties</i>	Post-Geneva I
1950	<i>United States Import Duties</i>	Post-Annecy
1952	<i>United States Import Duties</i>	Post-Torquay
1958	<i>United States Import Duties</i>	Post-Geneva II
1963	<i>Tariff Schedules of the U.S.</i>	Post-Dillon
1968–1972	<i>Tariff Schedules of the U.S.</i>	Post-Kennedy
1980–1987	<i>History of the Tariff Schedules of the U.S.</i>	Post-Tokyo

NOTE. The post-Kennedy and post-Tokyo documents contain MFN and column-2 rates of duty for the first year in the range, and staged MFN rates for the following years.

Some of our analysis (namely, in section 2.3) requires us to trace tariff rates on particular products over time. To assist in that analysis, we also digitized several concordances between 1930 and

1988, and relied on the digitized concordances of other authors to cover the period since 1988. To concord tariff lines in the *Foreign Commerce* (see table M.1) to their corresponding “paragraph numbers” used in the USID (1930–1963), we use the 1930 *Statistical Classification of Imports into the United States*.<sup>1</sup> Varieties tended to stay within tariff paragraphs over the USID period, so the concordance of goods from 1930–1963 is straightforward. To concord these paragraph numbers to their corresponding TSUS (1963–1988) numbers, we use of the crosswalk put together by the Tariff Commission in its 1960 *Tariff Classification Study*.<sup>2</sup> To concord items in the TSUS from 1963 through 1988, we use the *History of the Tariff Schedules* (see table M.1). To concord TSUS items to their HTS codes (1989–present), we use the concordance digitized by Feenstra (1996). Finally, to concord HTS items since 1989, we use the digitized concordance of Feenstra et al. (2002).

We end the description of the data with some summary statistics regarding legislated U.S. tariffs, shown in Figure M.1. The top panels, show the (simple) average ad-valorem and specific tariff on all dutiable items in each schedule. Despite not being trade-weighted, the patterns generally line up with trade-weighted average tariffs among dutiable items.<sup>3</sup> Our data do not cover the full breadth of non-dutiable imported items, though that is irrelevant for our analysis below.

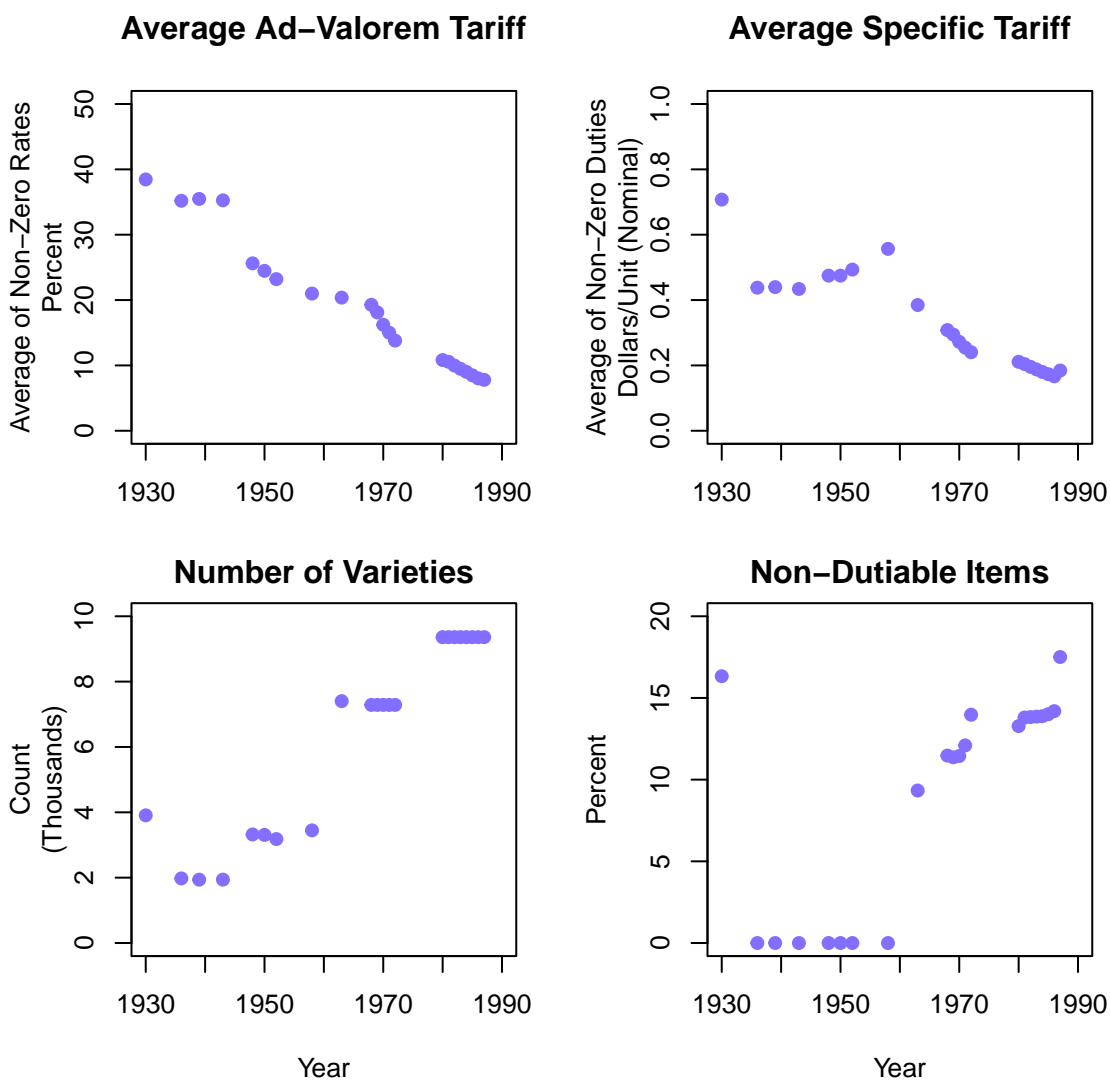
---

<sup>1</sup>To measure the full tariff schedule between 1930 and 1946, we carried the *Foreign Commerce* forward, updating it manually with the 1946 USID, guided by the *Statistical Classification*.

<sup>2</sup>Congress directed Tariff Commission (TC) to simplify tariff schedules and bring them up to date in the 1953 *Customs Simplification Act*. That process culminated in the 1960 report, which (after a few modifications) took effect with the 1962 *Tariff Classification Act*.

<sup>3</sup>See, e.g., [https://www.usitc.gov/documents/dataweb/ave\\_table\\_1891\\_2016.pdf](https://www.usitc.gov/documents/dataweb/ave_table_1891_2016.pdf).

Figure M.1: Legislated Tariffs: Summary Statistics



NOTE. This figure shows summary statistics for the tariff schedules that we use in our historical analysis. The top two panels show the (simple) average tariff rates among varieties with non-zero tariffs. The bottom-left panel shows the number of varieties in each schedule, and the bottom-right panel shows the percent of varieties that have a non-zero tariff.

## Appendix N: Transcripts, Minutes, ROPAs, MOAs, and MODs: The History of FOMC Communications

Understanding the logistical aspects of procedural transparency—which documents are released and when—is a necessary step in the assessment of procedural transparency in the “quality” sense. Since the inception of the modern-day FOMC, it has always communicated in some way the content of its meetings.

Table N.1: FOMC Publications: Banking Act of 1935 to Present (2015).

Date	Meeting Summaries	Detailed Accounts
1935–1967	Record of Policy Actions ( <i>Annually</i> )	Minutes ( <i>Confidential</i> )
1967–1975	Record of Policy Actions ( <i>90 Days</i> )	Memorandum of Discussion ( <i>5 Years</i> )
1975–1976	Record of Policy Actions ( <i>45 Days</i> )	Memorandum of Discussion ( <i>5 Years</i> )
1976–1993	Record of Policy Actions and Minutes of Actions ( <i>One Meeting</i> )	Transcripts ( <i>Confidential</i> )
1993–2005	Minutes ( <i>One Meeting</i> )	Transcripts ( <i>Five Years</i> )
2005–	Minutes ( <i>Three Weeks</i> )	Transcripts ( <i>Five Years</i> )

*Release lags in italics.* Adapted from Danker and Luecke (2005).

Table N.1 lists the various FOMC publications since 1935, along with their release lags.<sup>1</sup> While the nomenclature of the various documents has undergone several changes over the past 80 years, there have, in general, been two types of documents: detailed accounts of FOMC meetings and summaries. In general, the latter were more readily available to the public. Most of the changes in FOMC communications have occurred alongside calls for transparency, and they have formally come from Congressional pressure, legislation, and litigation. The exogeneity of these external pressures permits the study of how the FOMC’s procedural transparency responds to these calls. This section provides an overview of the logistical aspects of these responses.

<sup>1</sup>Statements, press conferences, and other releases are omitted, since the focus here is on documents whose primary purpose is to convey the meeting discussion.

The first significant step toward greater procedural transparency—in the sense of timeliness—came in response to the 1967 Freedom of Information Act (FOIA). Beginning with the April 1967 meeting, the Record of Policy Actions—a summary of the Committee’s policy actions and rationales—would be published after a 90-day lag (Danker and Luecke, 2005). And, for the first time, a transcript-like document—the Memorandum of Discussion—was to be released with a five-year lag. This set a precedent for publishing long accounts of FOMC meetings, but because the MOD was a heavily edited account, the 1967–1976 period is not included in the measurement of transparency reported below.<sup>2</sup>

The 1976 MODs were the last published by the FOMC; after five years of fighting a claimed FOIA violation, the Committee decided in 1981 to discontinue the MOD,<sup>3</sup> largely at the request of Chairman Burns (Lindsey, 2003). At this point, the Committee decided to release an expanded ROPA shortly after each subsequent meeting; effectively, a 30-day lag. At the time, the reason cited for the discontinuation of the MOD was that “the benefits derived from them did not justify their relatively high costs, particularly in light of the changes made in the [ROPA]” (ROPA, 1976–1993, 05/18/76). However, the more accurate reason seems to be “‘fear that Congress would request access’ [to the MOD] promptly” (Lindsey, 2003, p. 8) and, as an FOMC subcommittee indicated, “concern about the ability to conduct monetary policy, if the court required prompt release of the memoranda of discussion” (Meltzer, 2010, p. 976). The discontinuation of the MOD started a nearly 20-year period in which the FOMC published no detailed account of its meetings. Most FOMC members were aware that meetings were recorded, but they also believed that these tapes, used only for the production of minutes by Board staff, were recorded over after each meeting.

Contrary to what most members believed, Congressional inquiries (primarily headed by Congressman Henry González) and internal Fed investigations revealed that, in fact, these tapes had been maintained since 1976. In November 1993, the Committee agreed to publish all of the transcripts since 1976; by 1995 the decision was made to reinstate the publication of meeting tran-

---

<sup>2</sup>Though the similarity of minutes and MODs is shown above, this section of the time series should not be taken as measuring procedural transparency, due to the heavily edited nature of the MODs.

<sup>3</sup>See Goodfriend (1986) for a thorough account.

scripts after a five-year lag. In addition, the ROPA and MOA were now combined to form the “minutes.” In 2005, these minutes began to be released with a three-week lag.

All ROPAs, MOAs, minutes and transcripts were downloaded from <http://federalreserve.gov>, either in PDF format or plain text. Documents in PDF format were converted to plain text using optical character recognition (OCR) software.

## Appendix O: Definition and Relevance of Procedural Transparency

When used in the context of monetary policy, the word “transparency” can carry different connotations. To understand how the term is used here, table O.1 presents the five forms of transparency relevant to central banks, as defined by Geraats (2002).<sup>1</sup> Procedural transparency—my

Table O.1: The types of transparency relevant to central banks (Geraats, 2002, pp. F540)


<i>Type of Transparency and Description</i>
1. <i>Political Transparency</i> refers to openness about policy objectives and institutional arrangements that clarify the motives of monetary policy makers. This could include explicit inflation targets, central bank independence and contracts.
2. <i>Economic Transparency</i> focuses on the economic information that is used for monetary policy, including economic data, policy models and central bank forecasts.
3. <i>Procedural Transparency</i> describes the way monetary policy decisions are taken. This includes the monetary policy strategy and an account of policy deliberations, typically through minutes and voting records.
4. <i>Policy Transparency</i> means a prompt announcement and explanation of policy decisions, and an indication of likely future policy actions in the form of a policy inclination.
5. <i>Operational Transparency</i> concerns the implementation of monetary policy actions, including a discussion of control errors for the operating instrument and macroeconomic transmission disturbances.

focus in this paper—encompasses the procedure by which the accounts of FOMC decisions are released to the public via documents. What makes procedural transparency important is that increased procedural transparency presumably leads to increases in the other four types of transparency, and it is the mechanism through which the other four are manifested. For example, mandating that the Fed release the theoretical rule it uses to determine policy is a form of economic, policy, and political transparency. However, effective implementation of this policy hinges on effective procedural transparency because both the rule and deviations from it require detailed

<sup>1</sup>Geraats has written much about central bank transparency. See Geraats (2001), where these terms were first defined, or Geraats (2007) for other examples.

explanations. The same goes for the practice of establishing and explaining an inflation target—especially a “medium-term” target as is done in current practice. Simply put, the communications of the Fed are “a chance for [the FOMC] to say what [they] are up to and why” (Transcripts, 1976–2008, Alan Blinder, Jan. 1995).

Figure O.1: Google Trends: Fed Communications



The left panel presents results from the Google Trend queries “fed statement” and “fed minutes.” The y-axis represents the frequency with which a given phrase is searched on Google, and is normalized so that the highest frequency is 100. Thus, this graph does not say how these search terms rank among all other terms, but it does give information about when the terms are searched. Dashed lines are included at FOMC meeting dates. The right panel shows the Kullback-Leibler similarity, defined below, of the transcripts at each point in time with a query that contains the names of FOMC minutes and transcripts over time, and the word *communication*. The full list is “transcripts minutes record policy actions memorandum discussion communications.”

Fed communications also receive a considerable amount of attention from the public at large, and the FOMC itself. Figure O.1 supports this claim. The left panel is a graph from Google Trends—a service from Google that plots the “interest over time” of any search term—that shows how popular “fed statement” and “fed minutes” were over the year 2017. As expected, peaks in interest in the terms is in one-to-one correspondence with FOMC meetings; the three-week lagged release of minutes is also clearly noticeable. So, at the very least, there appears to be public interest in the content of FOMC documents.


The right panel of figure O.1 provides evidence that the FOMC discusses issues of procedural transparency in its meetings. Using the text-analysis techniques described below, it portrays the extent to which topics of procedural transparency were discussed at each FOMC meeting. This is

measured by computing the similarity between the transcripts and a list of procedural-transparency-related words (*transcripts, minutes, communication, etc.*). The similarity of this topic with the transcripts fluctuates meaningfully, coinciding with changes in communications policy—1976 marked the temporary end of transcript publication, and topics of procedural transparency persisted for a few years after that change. Since the early 1990s, changes in publication policy have been relatively frequent—in 1993, the FOMC decided to start making its transcripts public; and in 2005, the lag between the meeting and the release of the minutes was reduced. These changes are visible on the graph, indicating that a significant amount of discussion was behind each decision. Thus, given the role of procedural transparency in monetary policymaking generally, and the fact that both the public and the Fed pay close attention to the documents used here to measure it, the rest of this paper is devoted to its study.

## Appendix P: Number of Topics: Cross-Validation and Robustness

A common way to select topics for LDA is to estimate the model using a fraction of documents in the corpus, then compute how “perplexed” the model is by the held-out documents that were not used for estimation. In figure P.1, I show the results of performing this cross-validation for different values of  $K$  (from 10–100 by 10, and from 120–200 by 20). For each value of  $K$ , the sample is split into thirds—two thirds is used as a training sample, the term distributions are estimated, then using these distributions the topic distributions for the held-out documents are computed. The perplexity of these held-out documents is then computed. This is done five times for each value of  $K$ , shown by the black lines in the figure. The blue line shows the average of the black dots for each value of  $K$ . The red-dashed line shows the sum of squared of residuals from fitting two lines to the blue curve, where the two lines are split at a given value of  $K$ . This is called finding the “knee-point” of the blue line, or the point at which the perplexity drops off most-sharply. The knee of perplexity here is at  $K = 50$ , so 50 topics are used for the analysis in the paper.

Figure P.1: Cross-Validation Results



## Transparency Moving Average and Effectiveness




Figure Q.1: Estimates of  $\gamma$  for Different Moving Averages

## Appendix Q: Robustness for Moving Averages


Figure Q.1 shows the results of  $\gamma$  from equation (3.2), where the moving average of  $\bar{\tau}$  uses lag lengths from zero to 25. Stars are drawn whenever the level of statistical significance is at least 10%. All standardizations and scalings are performed as described in table 3.4. Notably, the coefficients are heavily skewed towards positive. Also, given how unlikely it is that the minute-transcript similarity for the most-recent meeting is observed (given that the transcript is not released for five years), the first few moving-averages, while negative, should be given little weight.

## **Appendix R: Tabular Representation of Figures 3.4 and 3.5**

	Baseline		TT State		Trans. Entropy		Time Trend		MPU	
	$\beta$	$\gamma$	$\beta$	$\gamma$	$\beta$	$\gamma$	$\beta$	$\gamma$	$\beta$	$\gamma$
Real Forwards										
2-Year	1.23 (0.27)	0.57 (0.15)	2.39 (1.36)	0.96 (0.27)	1.19 (0.26)	0.43 (0.18)	-2.00 (1.32)	0.43 (0.14)	1.28 (0.46)	0.59 (0.14)
3-Year	1.18 (0.24)	0.53 (0.13)	2.00 (1.16)	0.87 (0.24)	1.14 (0.24)	0.39 (0.18)	-1.72 (1.41)	0.40 (0.14)	1.48 (0.41)	0.53 (0.13)
5-Year	0.80 (0.22)	0.40 (0.12)	0.75 (0.88)	0.59 (0.20)	0.79 (0.21)	0.36 (0.20)	-0.81 (1.26)	0.29 (0.12)	1.33 (0.34)	0.39 (0.12)
7-Year	0.43 (0.21)	0.28 (0.12)	-0.22 (0.80)	0.39 (0.18)	0.45 (0.20)	0.37 (0.19)	-0.03 (0.95)	0.20 (0.10)	0.92 (0.32)	0.29 (0.12)
Nominal Forwards										
2-Year	1.49 (0.24)	0.56 (0.16)	2.34 (1.43)	0.94 (0.24)	1.45 (0.24)	0.50 (0.27)	-0.56 (0.93)	0.38 (0.12)	1.98 (0.46)	0.57 (0.14)
3-Year	1.22 (0.27)	0.53 (0.16)	2.35 (1.41)	0.95 (0.26)	1.19 (0.26)	0.47 (0.25)	-1.60 (1.39)	0.42 (0.15)	1.63 (0.48)	0.54 (0.15)
5-Year	0.63 (0.27)	0.40 (0.14)	1.13 (1.04)	0.70 (0.25)	0.63 (0.25)	0.39 (0.22)	-1.20 (1.59)	0.31 (0.15)	0.89 (0.38)	0.42 (0.15)
7-Year	0.24 (0.27)	0.27 (0.13)	-0.01 (0.89)	0.42 (0.24)	0.25 (0.23)	0.31 (0.23)	0.24 (1.34)	0.16 (0.13)	0.48 (0.33)	0.28 (0.15)
Real Yields										
2-Year	1.08 (0.27)	0.43 (0.16)	1.58 (1.16)	0.72 (0.25)	1.08 (0.25)	0.49 (0.20)	-0.57 (0.89)	0.30 (0.11)	1.16 (0.51)	0.46 (0.13)
3-Year	1.12 (0.25)	0.47 (0.14)	1.78 (1.16)	0.78 (0.24)	1.12 (0.24)	0.46 (0.19)	-1.00 (0.98)	0.34 (0.11)	1.24 (0.47)	0.49 (0.13)
5-Year	1.07 (0.22)	0.47 (0.13)	1.61 (1.05)	0.75 (0.22)	1.06 (0.21)	0.43 (0.17)	-1.10 (1.09)	0.34 (0.12)	1.32 (0.41)	0.48 (0.12)
7-Year	0.94 (0.21)	0.43 (0.12)	1.21 (0.92)	0.67 (0.20)	0.93 (0.20)	0.41 (0.17)	-0.88 (1.06)	0.31 (0.11)	1.26 (0.37)	0.43 (0.11)
Nominal Yields										
3-Month	0.16 (0.11)	-0.05 (0.11)	-0.72 (0.72)	-0.17 (0.14)	0.18 (0.11)	0.03 (0.12)	0.52 (0.21)	-0.04 (0.06)	0.53 (0.25)	-0.08 (0.08)
6-Month	0.40 (0.10)	-0.17 (0.08)	0.17 (0.26)	-0.28 (0.11)	0.40 (0.10)	-0.13 (0.10)	0.61 (0.24)	-0.10 (0.05)	0.52 (0.18)	-0.18 (0.07)
1-Year	1.00 (0.16)	0.19 (0.10)	1.00 (0.76)	0.27 (0.14)	1.00 (0.16)	0.26 (0.16)	1.00 (0.29)	0.08 (0.06)	1.00 (0.28)	0.21 (0.09)
2-Year	1.25 (0.19)	0.35 (0.12)	1.50 (1.00)	0.54 (0.17)	1.24 (0.18)	0.38 (0.20)	0.62 (0.42)	0.20 (0.08)	1.47 (0.35)	0.37 (0.11)
3-Year	1.29 (0.20)	0.42 (0.13)	1.80 (1.12)	0.68 (0.19)	1.27 (0.20)	0.41 (0.22)	0.02 (0.63)	0.27 (0.09)	1.59 (0.38)	0.43 (0.12)
5-Year	1.14 (0.22)	0.44 (0.13)	1.80 (1.11)	0.74 (0.21)	1.12 (0.21)	0.42 (0.22)	-0.63 (0.97)	0.31 (0.11)	1.45 (0.38)	0.45 (0.12)
7-Year	0.93 (0.22)	0.41 (0.13)	1.43 (0.98)	0.69 (0.20)	0.92 (0.21)	0.40 (0.21)	-0.58 (1.08)	0.29 (0.11)	1.22 (0.34)	0.42 (0.12)
10-Year	0.68 (0.21)	0.34 (0.12)	0.87 (0.84)	0.56 (0.20)	0.67 (0.20)	0.36 (0.20)	-0.10 (1.05)	0.23 (0.11)	0.96 (0.30)	0.36 (0.12)

Note: The table shows the coefficients  $\beta$  and  $\gamma$  estimated based on equations (3.2) (for columns labeled “Baseline”) and (3.3) (for the other columns). See the notes to figures 3.4 and 3.5 for details about the sample, column labels, and normalizations. Robust standard errors are in parentheses.

Figure S.1: Robustness of Transparency Measure to  $K$  and Sample Period

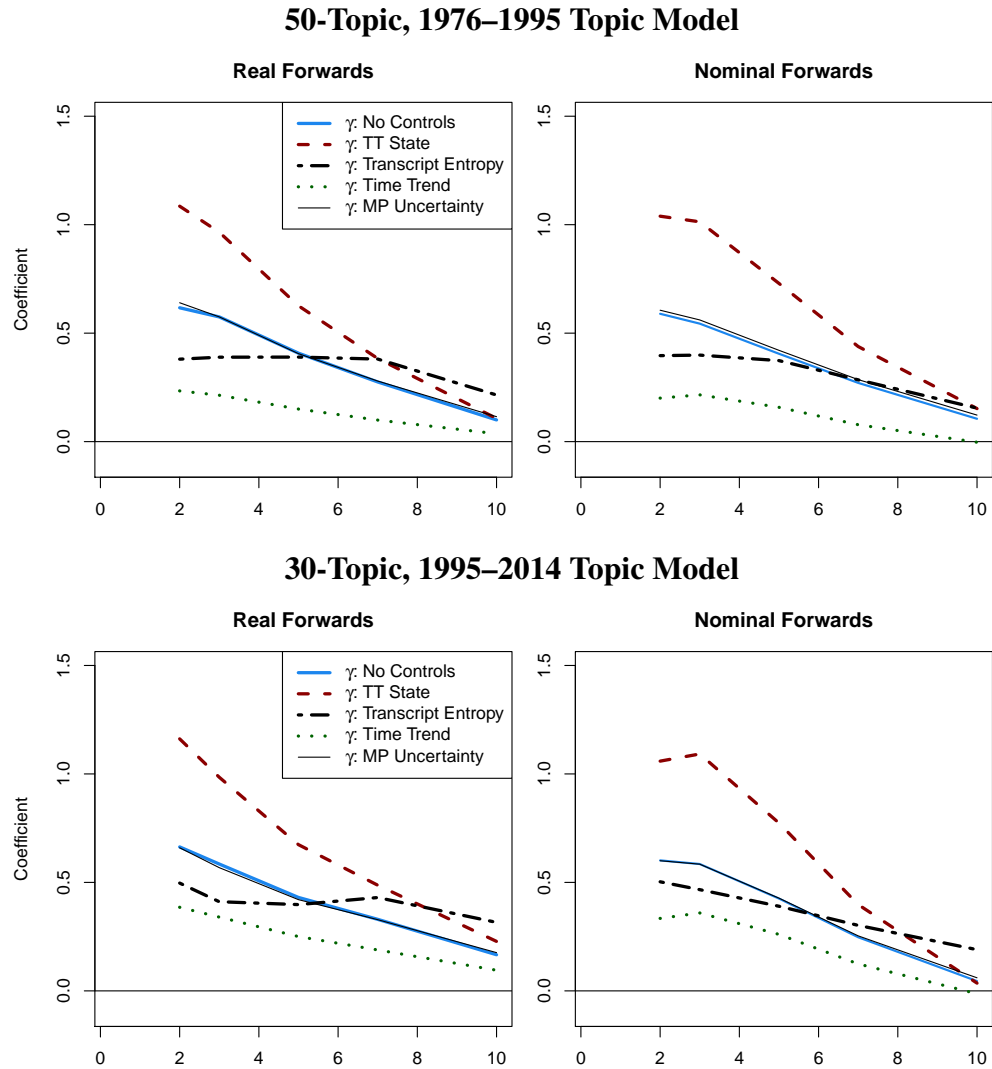


## Appendix S: Robustness to Text-Analysis Features

Figure shows the moving-average of the transparency measure using LDA topics estimated over the 1995+ sample (black-dashed line), which has a correlation of 0.86 with the full-sample measure. In addition, the gray-solid line shows the measure with  $K = 30$ , which has a correlation of 0.90 with the 1995+ sample measure with  $K = 50$ .

Figure S.2 shows the results of estimating (3.2) for interest rate forwards when the LDA model is estimated on the full sample

Figure S.2: Robustness of Regression Results to  $K$  and Sample Period



Note: The graphs show the coefficients  $\gamma$  estimated based on equation (3.3), where  $\bar{\tau}_{t-1}$  is estimated using the 1976–2014 LDA model (top panel) and the 30-topic 1995+ LDA model (bottom panel). Everything else is the same as in figure 3.5, so its note can be referenced for further details.