ESSAYS ON THE INTERFACE BETWEEN FINANCE AND TECHNOLOGY

Di Wu

A DISSERTATION

in

Finance

For the Graduate Group in Managerial Science and Applied Economics

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2016

Supervisor of Dissertation

Itay Goldstein, Joel S. Ehrenkranz Family Professor of Finance

Co-Supervisor of Dissertation

Amir Yaron, Robert Morris Professor of Banking and Finance

Graduate Group Chairperson

Eric Bradlow, K.P. Chao Professor of Marketing, Statistics, and Education

Dissertation Committee

Todd Gormley, Assistant Professor of Finance

Narasimhan Jegadeesh, Dean's Distinguished Chair of Finance

Nikolai Roussanov, Associate Professor of Finance

ESSAYS ON THE INTERFACE BETWEEN FINANCE AND TECHNOLOGY

ACKNOWLEDGEMENT

I am infinitely grateful for the help and guidance from my dissertation committee members: Itay Goldstein (Co-Chair), Amir Yaron (Co-Chair), Todd Gormley, Narasimhan Jegadeesh, and Nick Roussanov. I would also like to thank Kenneth Ahern, Ravi Anupindi, Simon Gervais, Gerard Hoberg, Rajkamal Iyer, Donald Keim, Bryan Kelly, Naveen Khanna, Hyunseob Kim, Leonid Kogan, Timothy Loughran, Bill Mcdonald, Jun Pan, Jonathan Parker, Uday Rajan, Ivan Shaliastovich, Rob Stambaugh, Alireza Tahbaz-Salehi, Luke Taylor, Paul Tetlock, Thomas Sargent, Philip Strahan, Michael Weisbach, Lu Zhang, Da Zhi, my Ph.D. colleagues and other Wharton faculty, conference participants at the WFA-CFAR Conference, the American Finance Association Annual Meetings, Chicago Quantitative Alliance Spring Meeting, and seminar participants at the Federal Reserve Bank of Atlanta, Wharton School of the University of Pennsylvania, Texas A&M University (Mays), Boston Collage (Carroll), University of Chicago (Booth), Emory University (Goizueta), Massachusetts Institute of Technology (Sloan), University of Southern California (Marshall), University of Notre Dame (Mendoza), Michigan State University (Broad), Ohio State University (Fisher), University of Connecticut, Washington University in St. Louis (Olin), and University of Michigan (Ross), for helpful comments and suggestions. I would like to thank my wife, Jun Li, for her countless times of help and encouragement. This dissertation is supported by the National Science Foundation (NSF) Award #1547987, a Research Grant from the Rodney L. White Center for Financial Research, and the BlackRock Applied Research Award.

ABSTRACT

ESSAYS ON THE INTERFACE BETWEEN FINANCE AND TECHNOLOGY

Di Wu

Itay Goldstein

Amir Yaron

In the first chapter "Shock Spillover and Financial Response in Supply Chain Networks: Evidence from Firm-Level Data", using machine learning methods on firm-level textual disclosures, I construct a large-scale dataset featuring firm-specific shocks to production. I map these shocks into a unique, hand-built network of firm-level supply chain connections to empirically quantify how these localized shocks affect remote firms along the chains. Surprisingly, contrary to prediction by typical network theories, these firm-specific shocks impact the revenue of firms even up to 4 connections away from the origins. This pronounced spillover effect is explained by three features–uneven distribution of monopolistic power, variations in supplier substitutability, and different inventory levels–that are salient in the data but usually not present in existing models. In addition, firms seem to respond to these spillovers by increasing their working capital and financial leverage. Moreover, the stock market reacts to shock spillovers from distant connections with slower speeds: post-shock abnormal returns are persistently negative for up to 40 days.

In the second chapter "Deciphering Fedspeak: The Information Content of FOMC Meetings", we present a new approach to quantify the economic and policy content of Federal Reserve communications by dissecting the Federal Open Market Committee (FOMC) meeting minutes into distinct economic topics, and simultaneously extract the tone and uncertainty level of each topic. We use market reaction to objectively assess the relative informativeness of each topic, and we find significant incremental informational value from the topic contents, despite that the minutes are released several weeks after the original meetings.

Furthermore, we find evidence consistent of the Fed possessing superior information, which is then transmitted to the market through the language of the minutes.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF ILLUSTRATIONS

CHAPTER 1 : Shock Spillover and Financial Response in Supply Chain Networks: Evidence from Firm-Level Data

1.1. Introduction

The most prevalent type of inter-firm relationships is the physical network of supply chains. Its large scale and complex, changing structure is evident from Figure 1.1, which plots a slice of the network between the 400 largest firms in the tech industry in 2002 and 2015. With globalization, supply chains have become even longer, denser, and more clustered. The additional verticalities introduced by these deeper firm-level interconnections remains unexplored, as typical studies on spillovers and externalities in finance focus on the *existence* of externalities and their effects on firms' *direct* connections. By contrast, we have limited information on the *extent* of these externalities, measured by how deep they penetrate along the linkages to *remote* connections beyond the direct linkages.

This paper directly quantifies the extent of these externality spillovers: I hand-build two firm-level datasets to empirically examine the economic effect of firm-specific production shocks on the origin firms' remote customers along the supply chain, beyond their direct, tier-1 connections. I then explore the implications of this remote spillover for the stock prices and corporate policies of firms further down the chains.

While there are many recent theoretical papers on the aggregation of shocks through production interconnections such as Acemoglu et al. (2012), Gabaix (2011), and Kelly et al. (2013), little empirical evidence exists on their firm-to-firm implications due to the lack of well identified data on firm-specific shocks, and the lack of granular data on firm-level networks of supply relationships.[1] This paper contributes two new datasets to address these limitations in identification and measurement. Specifically, I first build a 20-year database of over 8,000 firm-specific production shocks of a large variety of types and firms. I identify

[1] Existing data is either sector level, such as the make-use tables aggregated by the Bureau of Economic Analysis (BEA), or only capture a small fraction of relationships between small suppliers and very large customers, such as the Compustat Segments database.

these shocks through a topic-based textual analysis of over 5 million individual firm disclosures. This allows for well-identified, firm-level empirical examinations. I then map these shocks into a 20-year network of over 1 million supply chain relationships between publicly traded firms globally. I construct this firm-level supply chain network from the disclosure data and three other public and proprietary data sources. Tracing the exact origin of each shock to the network, I directly examine if these firm-specific, idiosyncratic shocks spill over along the linkages. I measure their impact on the economic outcomes, corporate policies, and stock prices of origin firms, their closest connections, and remote firms beyond the first-tier connections.

I find several new results. First, although typical theories on production networks predict that firm-specific shocks quickly decay after the first link, I document that, surprisingly, these shocks cause substantial impact to firms even up to 4 connections away from the origin: on average, an idiosyncratic shock that causes 1% decline in revenue growth for the origin firms causes 0.82% decline for their closest connections, but also causes 1.03%, 0.87%, and 0.40% decline for 2nd, 3rd, and 4th-tier connections.

I then demonstrate that the data exhibit three salient features, which are absent from typical theories, yet are particularly conducive to such prolonged and pronounced shock spillovers. First, (monopolistic) market powers are unevenly distributed along the supply chains. Firms located further away from origins tend to have lower market power, thus are less able to change prices to their customers, and at the same time face higher-powered suppliers who are more able to raise prices and pass more impact onto them. This dual price-quantity effect causes larger impact and longer spillovers than predicted by models with perfectly competitive firms. Second, the substitutability of suppliers and third, the level of inventories, also vary along the supply chains. These factors can also significantly affect the magnitude of spillovers. I first establish the validity of these channels through firm-level tests, where I show that they are indeed significantly related to the impact of spilled-over shocks. I then provide the network-level evidence for the market power channel by documenting that firms further down the supply chains have lower average market power.

Next, I develop a series of robustness checks to solidify the results' internal and external validities. First, although the distinguishing feature of the shock data is that they are derived from firm-specific disclosures, one still needs to verify that they do not have systematic *causes* (i.e. caused by macroeconomic fluctuations), nor systematic *effects* (i.e. their initial impact hits only the origin firm). During data construction, the shock disclosures are classified by a Bayesian topic model called Latent Dirichlet Analysis (LDA). This results in multiple, distinct types of causes, from natural disasters (earthquakes, floods), to production glitches (machine breakdowns), etc. I examine the impact of each category at a time, as well as a sub-category consisting of only localized plant fires, and find similar results across categories. This confirms that the sample is not contaminated by shocks with potentially systematic effects and that most shocks from my data are indeed localized. Second, firms might only report the largest shocks, leading to overstatement in coefficient estimates. I use the exogenous enactment of a much stricter SOX-Act provision to rule this case out. Third, I conduct a series of falsification tests to ensure that the shocks are indeed correctly mapped to the network data. A host of tests for other issues such as strategic reporting, the role of private firms, etc are discussed in the text.

Having documented the main economic facts, I then explore the stock market implications of these long spillovers. I first document that prices are quickly adjusted for firms with their own shocks or shocks from a directly connected supplier, with a two-day abnormal return of -3.98%. However, if the shock is originated from further up the supply chains, then the market response, while still big in magnitude, is much slower in speed: there is a persistent drift in returns for up to 40 days. Exploiting variations of my shock data where some shocks are disclosed by firms themselves while others are disclosed by their indirect connections, I demonstrate that information processing constraints, perhaps similar to those uncovered by Cohen and Lou (2012), are likely responsible for this slow reaction, giving rise to potentially profitable trading opportunities for investors more adept at analyzing the complex structure of supply chains.

My final set of tests explores the responses from corporate policies and capital structure to shock spillovers. To do so, I first examine the languages of 10-K/Qs from the same firms that disclose shocks and uncover a series of policies that firms *say* they would adjust. I then empirically test whether firms actually *do* adjust these policies and the intensity of such adjustments. My evidence suggests that, after shock spillovers, firms increase their inventories, cash holdings and capital expenditures. I then show that this response is much larger to shocks originated from more closely connected sources, than from more remote firms up the supply chains.

This paper contributes to the literature on four fronts. First, a burgeoning literature in finance explores spillovers and externalities along various types of interconnections related to personnel (Shue (2013), Cohen et al. (2008), Cohen et al. (2010), Maturana and Nickerson (2015)), geography (Poll et al. (2015)), trading, financing and insurance (Cohen-Cole et al. (2014), Bajo et al. (2015), Billio et al. (2012), Afonso et al. (2015)), product (Hoberg and Phillips (2015), Rauh and Sufi (2012), Foucault and Fresard (2014)), supply chain (Barrot and Sauvagnat (2014), Titman and Wessels (1988), Banerjee et al. (2008), Fee et al. (2006), Hennessy and Livdan (2009)), or internal connections within conglomerates (Schoar (2002)), and the effect of these externalities on firm behaviors. A large body of literature in operation management, such as Anupindi and Akella (1993), Tomlin (2006), and Ang et al. (2014), focus on how to mitigate these externalities operationally. However, little is known about the *extent* of these externalities in terms of how *deep* their effect can penetrate along connections. This paper builds upon this literature by directly quantifying the extent of these externalities according to how far they propagate vertically down the supply chain linkages. I demonstrate that, at the micro level, these externalities spill over much deeper beyond just one firm, with their effect reaching up to the 4th connection in some cases. Therefore, at the macro level, the aggregate effects are potentially even larger than existing studies suggest.

This paper also contributes to the literature on the identification and measurement of firm-specific shocks and their propagation. Existing studies in finance and operations management have examined shocks in the context of bankruptcies (Hertzel et al. (2008), Kolay and Lemmon (2011)), financial distress (Hortaçsu et al. (2013)), natural disasters and temperature variations (Barrot and Sauvagnat (2014), Boehm et al. (2014), and Bergman et al. (2015)), and supply disruptions (Hendricks and Singhal (2005)). Moreover, Leary and Roberts (2014) identify shocks by extracting the idiosyncratic components of equity returns using asset pricing models. Giovanni et al. (2014) employ a similar approach in extracting the idiosyncratic component of firm sales. Other papers, such as Strebulaev and Whited (2011) and Gourio (2008), recover such shocks with structural models on firm production and investment. However, so far we have limited information on the exact nature of those shocks. This paper builds upon this literature with a different source of data. With the actual texts of firm disclosures and news, I directly observe the actual events behind the shocks. This direct way of capturing the source of idiosyncratic production shocks results in additional granularities that are helpful not only for identification, but also for the precise measurement of economic magnitudes of a wide variety of shocks in other contexts as well (e.g. financing shocks or systematic shocks).

Moreover, this paper complements existing work on asset pricing and shock aggregation in production networks. Compared to sector-level studies on aggregate volatilities and production networks such as Ahern (2012) and Foerster et al. (2011), this paper directly captures firm-specific production shocks, thus enriches the firm-level microfoundations of shock aggregation. This paper complements the work on correlated supplier-customer stock returns, such as Cohen and Frazzini (2008) and Menzly and Ozbas (2010), by measuring the response of stock returns directly to firm-specific production shocks.

Finally, the empirical results on the link between market power and shock spillovers provide motivation for future theory developments on production networks that incorporate richer sets of frictions such as monopolistic competition, and the asset pricing relationship between the spillover effects and network-related systematic risks and expected stock returns. I explore some of these issues in related works.

The rest of the paper is organized as follows. Section 1.2 describes my sample and data sources. Section 1.3 reports the empirical results on shock spillover in the supply chain network, and their effect on firms' revenue, cash flows, and gross margins. It also isolates several factors that could give rise to, and affect the magnitude of, these spillovers. A series of key robustness tests immediately follow in Section 1.4 to ensure the validity of these spillover results. Then, Section 1.5 discusses the responses, from both managers and the stock market, to these spillovers. Section 1.6 concludes.

1.2. Data and Identification

I derive the main empirical inferences in this paper by mapping a hand-collected, 20-year dataset of firm-specific idiosyncratic shocks into a hand-built dataset of firm-to-firm supply relations between publicly traded firms globally. This section describes my methodology to construct these datasets. A more in-depth review of the data construction methodology can be found in Appendix A.1.2.[2]

1.2.1. Firm-Specific Idiosyncratic Shocks

Data Sources and Initial Sample Selection

Using a two-step textual analysis technique, I extract data on firm-level idiosyncratic shocks from the language of corporate disclosures and news. The first step extracts the list of *supply shocks* from these texts, and the second step isolates the *firm specific, idiosyncratic* ones from the first list. For the first step, I download and process the texts produced between 1994 and 2015 from the following sources:

1. SEC Form 8-Ks filed in the EDGAR system: In addition to regularly scheduled disclosures such as 10-K/Qs, public companies in the United States are required to report

[2] Please contact the author if interested in using the described data for research purposes.

a wide variety of material corporate events on a more timely basis, in the form of 8-K filings, or "current reports." Such disclosures were optional but became mandatory following the passage of the Sarbanes-Oxley Act. Correspondingly, many firms disclose a wide variety of supply chain-related events and disruptions in their 8-Ks. Additionally, exact event dates are included in HTML headers accompanying the main filings. I parse these data from the headers and merge them with the main texts to create one unique date-stamped, item-coded text string per 8-K filing. The filing firm's identity is included in the SEC's Central Index Key (CIK), which I merge with GVKEY from Compustat.

2. Press releases filed through the Dow Jones Newswire: I obtain the main texts of each release and match the name of the disclosing firm back to Compustat.

3. Company-specific news from Capital IQ: this includes most firm-related news in a time-stamped, machine readable format.

For each text document in the above list, Step 1 of the textual analysis uses a set of keyword filters that captures events containing keywords from the following groups:

- Supply chains: *supplier, supply chain, shipment, raw material*, etc;
- Events: *disruption, shortage, delay*, etc;
- Shocks: *unexpected, sudden, shocks*, etc;

Specific keyword filters can be found in Appendix A.1.2. This step results in a set of 24,838 events corresponding to supply-related shocks. For 19,771 of such shocks, I can identify the exact origin of the shock by iteratively searching the texts for the list of all known firm names extracted from Bloomberg. Of these shocks, 14,043 (71.03%) are shocks to the disclosing firm itself and 5,728 (28.97%) are disclosed as a supply shock from a named supplier firm. In the former case, the origin of the shock is identified as the disclosing firm itself. This distinction proves useful for a key robustness check in Section 1.4.1.

Classification of Shocks by Type

The second step of the textual analysis represents a key innovation of my data construction: I fit a Bayesian *topic model* from the Latent Dirichlet Allocation (LDA) family on the output disclosure collection from the previous step. The topic model is designed to infer the type, or nature, of each disclosed event. This paper is the first in finance to employ the LDA model to classify the nature of firm-level shocks from disclosures.[3] In this subsection, I briefly review the reasoning for the model, while deferring the exact model specification and computation steps to Appendix A.1.2 for interested readers.

In short, the LDA model assumes a simple, two-distribution data generating process where each disclosure is generated from a (latent) distribution over a collection of topics (i.e. shock types), each of which is, in turn, a distribution over the words in the English vocabulary. For example, a document that discusses the impact of a plant fire of its supplier should be represented by a topic distribution that places high weight on a topic that places high weight on words such as *fire* and *flame*. By contrast, a topic that places high weight on *earthquake* and *flood* should receive a low weight in this distribution.

However, the two distributions are unobservable from the point of the researcher. The advantage of probabilistic topic models is that, using Bayesian techniques, such models efficiently infer the hidden distributional properties from the observable data (i.e. the collection of documents). LDA represents one particular parameterization of the model: I first assume that these two latent distributions belong to the Dirichlet family. Then, armed with this functional form and the observed words in each disclosure, I compute the posterior (i.e. empirical) topic and word distributions using the standard Bayes Theorem. These empirical distributions are the main outputs of the model. The only inputs in LDA are the document texts and the number of topics. As such, compared to a manual classification approach, researcher-induced subjectivity and bias are minimized, and a much larger number of disclosures can be classified within a short period of time. Intuitively, the model would achieve

[3]For a list of LDA applications and an evaluation of their effectiveness, see Blei et al. (2003).

a satisfactory performance if the top words representing each topic are distinct from each other.

I fit the LDA algorithm with 20 topics on the collection of 19,771 disclosures.[4] The first outputs are the word distributions that identify the topics, and the top 5 keywords for each topic are reported in Table 1.2 from Appendix A.1.2. Here the top words are very distinct while intuitively and clearly identifying each topic. From these top keywords alone, an average human reader would not have any confusion about the nature of each topic.[5] The second output is the topic mixture for each disclosure, i.e. the proportion of each disclosure devoted to each topic. I keep the disclosure only if it exhibits more than 95% of a single topic. This results in a collection of 12,337 documents, each classified into one of the 20 topics.

Identifying Idiosyncratic Shocks

From the 20 types, I then isolate the firm-specific, idiosyncratic ones. First however, note that the meaning of "idiosyncratic shocks" differs in different research contexts in finance. In my setting of shock spillovers, the following definition applies:

Definition 1 *A shock is firm-specific and idiosyncratic if it is 1) not caused by exposure to common, systematic factors, and 2) its immediate effect is clearly limited to the origin firm only.*

In other words, a firm-specific shock is clearly identified from a disclosure if the disclosed event does not have systematic *causes* (i.e. caused by macroeconomic fluctuations), nor does it have systematic *effects* of its own (i.e. its initial impact hits only the origin firm). Both statements serve as exclusion conditions to insure the validity of the central identifying

[4]The main algorithm and the Gibbs sampling programs are implemented in the C++ programming language. The classification remains robust from as few as 15 topics to as many as 45 topics, after which overlapping redundancies appear. See Appendix A.1.2 for details. For research purposes, interested readers can also contact the author for programs used in the analyses.

[5]Jegadeesh and Wu (2015) ask a team of human readers to manually classify a list of FOMC meeting minutes. They agree with the LDA model for the vast majority of cases.

assumption of this paper: any effect that a shock might cause to other firms is attributable to the spillover effect through inter-firm linkages, and not due to original effect of the shock itself. The first condition (idiosyncratic cause) is easier to verify, and I discuss that below. The second, more nuanced condition (idiosyncratic effect) requires more extensive tests and is thus devoted its own Section 1.4.1.

I first focus on the causes of these shocks. Because the LDA provides an intuitive classification of shock types, the first condition, that the shocks are not caused by systematic factors, can be validated by examining the nature of these types. Specifically, the 20 types can be grouped into three levels of "idiosyncrasy": First, 6 types of events are related to macroeconomic conditions and industry-wide factors, e.g. parts shortages caused by unexpectedly high demand from other firms or consumers. Second, 6 types are related to events with more granular causes, but it might still be possible to relate these types to industry-wide issues. For example, events related to labor strikes might well be caused by firm-specific factors, but one still cannot completely rule out the influence of industry-wide labor unions that can possibly coordinate these strikes with both origin firms and their connected customers. Finally, the last level contains 8 shock types, each one of which can be attributed to idiosyncratic causes with little ambiguity. Such events include various natural disasters, manmade disasters such as fire or crime, production glitches such as power outages and unexpected machinery breakdowns, and technology adoption failures such as IT glitches and attacks, etc. I include several concrete textual examples for each category in Appendix A.1.3.

Because events in the last group are clearly not caused by systematic factors, they are the best candidates for well-identified firm-specific shocks. I therefore keep only these events in my sample and group these 8 types into five major groups.[6] The resulting sample is summarized in Table 1.1 below. Overall, I identify 11,191 idiosyncratic shock events, covering 2,193 unique firms or 20.06% of the sample firms. At the firm level, they are

[6] In untabulated results, I also include the second group e.g. labor strikes. The results are slightly stronger both economically and statistically.

clearly of very low frequency, with each firm on average having 5 shocks.

[Insert Table 1.1 and Figure 1.2 here]

Figure 1.2 plots the total number of these idiosyncratic shocks over time, and the distribution by major types. Except for the first three years in the sample, the number of shocks does not exhibit any particular trend or correlation with systematic factors such as business cycles, further confirming that the shocks are indeed caused by idiosyncratic issues. A potential concern is that the number of natural disaster-related shocks seems to increase during the years of 2005 and 2011, the years when hurricane Katrina and the Japan earthquake took place. I address this and other concerns related to the shocks' *effects* in Section 1.4.1.

1.2.2. Firm-Level Supply Chain Network

In order to empirically examine the spillover effect of the firm-specific shocks described in the previous subsection, I construct a comprehensive, firm-level network of supplier-customer relations between publicly traded firms globally, with historical coverage beginning at 1994. Each shock in the previous dataset can then be mapped precisely to this network. Specifically, for each year t, I capture the list of all supply relations between each supplier j and customer i, inferred from the following sources:

1. The same collection of firm disclosures described in the previous subsection. In addition to operations-related disclosures, the SEC also requires firms to disclose the *formation and termination* of important business relationships, a large portion of which are supply chain relations. I use a robust name-matching algorithm to extract the identities of both parties and record formation and termination dates from the collection of 8-K filings, press releases, and firm-specific news. This procedure produces 414,355 supply relations among firms.

2. The Supply Chain Analytics databases from Bloomberg and Revere Data Systems,

both of which conduct public and proprietary research that identifies supply chain relationships between publicly traded firms globally. I extract these data with proprietary APIs provided by these sources and identify 1,190,474 relations.

3. Shipment-level data from US Customs Bill of Lading and a leading business casualty insurance company. Both datasets provide a comprehensive account of all import/export goods that clear the US Customs at the ports of departure/arrival. Detailed identities of both supplier(shipper) and customer(consignee) are also recorded. I extract these information using a similar name-matching algorithm that results in 652,932 relations.

The resulting dataset contains 2,257,761 relations covering 23,059 publicly traded firms. I use the following criteria to construct my final sample of firms:

- My dataset identifies firms using the International Securities Identification Number (ISIN). I match the ISIN with GVKEY from the Compustat Global database to retrieve accounting data. I exclude all firms for which I am either 1) not able to match ISIN to GVKEY's or 2) otherwise not able to obtain accounting data from Bloomberg.

- At minimum, my tests also use market capitalization and return on total assets as control variables. I exclude all firms for which these data are unavailable.

- The production process for financial and other services are likely to be different from that of other goods. I exclude all financial and personal services firms (SIC code 6000-7999) from the sample.

Overall, the final sample captures 10,930 unique firms and 1,007,998 relationships from 1994 to 2015. Table 1.3 provides summary statistics of sample firms and supply relationships captured by this sample. The mean market value is $2.045 billion and the book-to-market ratio has a mean value of 0.715. The network data identifies over 90 links per firm over the

1994-2015 period, and each link on average persists for 7 years. In addition, Panel C of Table 1.3 reports the average statistics for firms located on the path of the shocks (i.e. connected at a distance of $n = 1, \ldots, 4$ from the firms where the shocks originate). On average, firms located further downstream (higher number of connections) from the shocks are slightly smaller and younger in age. Except for that, firms located at different network positions are similar in most other aspects such as book-to-market ratio, P/E ratio, leverage, etc.

[Insert Table 1.3 here]

For a smaller set of relationships (473,759), I am also able to capture the value of the relationship as a fraction of the customers' cost of goods sold (COGS).The summary statistics for these firms are very similar to the overall sample and are reported in Table 1.3 above. For this sample, the average total COGS explained by the supply shares is 34.04%. I use this smaller sample in a series of robustness tests, discussed in Section 1.4.

1.2.3. Other Relevant Data

All quarterly accounting data are from Compustat Global and Bloomberg. I obtain daily stock return data primarily from CRSP and compute the same for international firms using daily closing prices from Compustat, manually adjusted for stock splits and dividends. I use these data to construct firm-level outcome and control variables, which are discussed as they appear in subsequent texts. For interested readers, I also summarize the list of variables used in this paper and their construction methodology in Appendix A.1.1.

1.3. Spillover of Firm-Specific Idiosyncratic Shocks in Network

This section provides empirical evidence on the substantial spillover of firm-specific shocks along the firm-to-firm interconnections in the supply chain network, and the effect of such spillover on the revenue and cash flows of firms that are 1) initially hit by the shocks, 2) closely connected to the shocks' origins, and 3) remotely connected to the origin beyond the first-tier connections.

1.3.1. Hypothesis Development

Conceptually, the shocks captured by my disclosure data represent firm-specific innovations in their respective productivity processes. First, in models without networks, such shocks do not impact any firm beyond the origin. Second, in most theoretical models with production networks, such innovations are interlinked by the network connections (represented by the network's adjacency matrix). In the aggregate, these interlinkages introduce additional covariances among firms' outputs, thus changing aggregate output and consumption volatilities. However, it is easy to derive that along each chain of links, e.g. S→C1→C2→ ..., the effect of a shock quickly decays after each link.[7] This scenario, which serves as the null hypothesis, is that a shock to S would not significantly affect C1 and other firms beyond it.

So far, there is no empirical evidence on how far these shocks travel after C1 and whether they do quickly dissipate. However, a large amount of disclosures from my text sample do suggest that the effect of these shocks might persist beyond C2-type firms and even significantly affect more remotely connected firms. Figure 1.3 plots the exact timeline of such an example. Here, the factory of a hard drive component supplier (S:Nidec) was damaged during a flood in 2011, causing production disruption and leading to lower revenue growth. Its immediately connected customer (C1:Seagate) disclosed this disruption almost immediately, and verified that its own production facilities were not affected. However, due to input shortages, its output was also lower, while at the same time, due to its large market share in the hard drive industry, it was able to charge a higher price for these outputs, partially offsetting the quantity drop and leading to a more moderate decline in revenue growth.

Surprisingly, C2 firms on the supply chain, Dell and HP, both experienced larger revenue

[7] To see this, consider the Cobb-Douglas production setting in Acemoglu et al. (2015) with output Y_i, labor input L_i and intermediate good inputs X_{ji}: $Y_i = A_i L_i^{1-\alpha} \left(\prod_{j=1}^{N} X_{ji}^{\gamma_{ji}} \right)^{\alpha}$. Without an asymmetric network (e.g. if we only have a single chain), if $\alpha < 1$, i.e. labor has a non-zero share in total factor inputs, then at each subsequent connection k, the impact of a shock to firm j would decrease by a factor of $\alpha\gamma_{kj}$ as it travels downstream along the chain. With average $\alpha = 0.60$ and $\gamma_{kj} = 0.30$, the effect of the shock is 82% lower after each link.

declines, and both attributed them to supply chain disruptions caused by the flood. Particularly, they mention "significant and immediate" increases in supplier prices, coupled with their inability to pass this increase to customers (due to the more competitive nature of the computer manufacturing business compared to hard drive manufacturers), as the reason for these lowered revenue and earnings numbers.

This example suggests that, contrary to predictions by existing theories, which assume smooth production functions and competitive firms, significant spillover of firm-specific shocks to remote connections can occur if firms along the chains are not perfectly competitive and differ in characteristics such as market power. Consider a supply chain consisting of three firms, S$\rightarrow$C1$\rightarrow$C2, where S experiences a shock causing lower outputs to both itself and C1. Existing network models predict that the shock would not significantly effect C2 if C1 is perfectly competitive. However, suppose both C1 and C2 have monopolistic competitive powers. In this case, because C2's input is now more scarce, C1 would raise its price above the competitive level in order to recoup some losses and pass some effect to C2. This behavior is extensively discussed in firm disclosures and illustrated in Figure 1.3. If C1 has a high market power, it would be able to raise prices significantly (Seagate, the C1 firm in the previous example, raised prices by 20% shortly after the flood). If C2 has a low market power, it cannot raise prices as effectively while simultaneously facing a high-powered supplier. In this case, C2's output will be even further below the competitive level predicted by existing models, thereby further extending the spillover effect of the original shock.

In essence, at a granular level, significant deviations from perfect competition could lead to significant shock spillover effects beyond the prediction of existing models that assume perfect competition at all positions. In addition, Section 1.3.3 empirically examines two auxiliary factors that might also contribute to shock spillovers: input substitutability and inventories. Given the above framework, I test the following alternative hypothesis against the theory-predicted null of no spillovers:

Hypothesis 1 *Given additional frictions such market power, input substitutability and inventory, firm-specific idiosyncratic shocks could spill over significantly beyond both the origin and its closest connections, and significantly impact the economic outcome of a larger number of distant connections.*

I test the this hypothesis first using the following regression, which measures average outcomes across all shocks:

$$Y_{it,t+k} = a + \sum_{n=0}^{10} b_n D^n_{i,t} + cX_{i,t} + F_{i,t} + \epsilon_{i,t}, \tag{1.1}$$

where $D^n_{i,t}$ is a dummy variable that equals to 1 if one of firm i's suppliers from a distance of n connections (up to 10) experiences an idiosyncratic shock in quarter t. $Y_{it,t+k}$ is the k-quarter growth rate in revenue, operating cash flow, or change in gross margins. $X_{i,t-1}$ is a vector of lagged controls including market capitalization ($Size$), book-to-market ratio (BM), P/E ratio (PE), leverage ratio (Lev), return on assets (ROA), and inventory (INV), all defined in Appendix A.1.1. $F_{i,t}$ is the set of fixed effects including: industry×year, fiscal quarter, and state/country fixed effects. The coefficients of interest is b_n, which measures the average difference in revenue and other growth rates between firms hit with a shock spilled over from a distance of n connections, and firms never hit with any shocks. Hypothesis 1 predicts that the estimates for b_n will be significantly negative, not only for $n > 0$, but also for many of $n > 1$ as well.

1.3.2. Spillover Results

I first fit Regression (1.1) using 4-quarter revenue and operating cash flow growth rates,[8] as well as changes in gross margin, and report the coefficient estimates for $n = 0, \ldots, 4$ in Table 1.4. The coefficient b_0 measures the effect of the shock on the originating firm itself, and none of the coefficients b_n are statistically significant at the 10% level when $n \geq 5$.

[8] I also use revenue growth rates from t to $t+1$, $t+2$, and $t+3$ quarters. For most distances, the impact is concentrated around the first quarter i.e $(t, t+1)$.

[Insert Table 1.4 here]

The first column of this table reports the impact of the shocks on originating firms themselves. The coefficient estimate for b_0 is -0.0258, compared to a sample mean revenue growth rate of 0.1061 with a standard deviation of 0.2885. Therefore, on average, the shocks captured by my disclosure data indeed have significant impact on the firms that are directly hit by them, causing revenue growth to be about 24.3% lower than average or 8.9% standard deviation lower.

The next four columns document the significant spillover of shocks to as far as 4 connections away from origin. Starting from the closest connection (n=1), these firm-specific shocks cause substantial impacts to the revenue growth of customer firms in the next 4 tiers of connections. Following the shock, revenue growth for the closest customers is about 21.6% below the mean. This is smaller in magnitude, but similar in statistical significance, to Barrot and Sauvagnat (2014), who, using hurricane and Compustat segments data, find that the largest customers (> 10% revenue) of firms whose headquarters are located in the disaster area subsequently experience revenue growth up to 35% lower.

Surprisingly, the effect of my firm-specific shocks continues to spill over to the customers of these customers (n=2), and their customers (n=3), and again to their customers (n=4). The average effect ranges between 11.7% to 35.5% lower than the mean. While the difference between these coefficients are not statistically significant at the 10% level, the coefficients themselves are all significantly negative, with t-statistics below -3.00. This provides strong evidence of the existence of frictions e.g. market power, that are not present in existing theoretical models. This results in the significant spillover of firm-specific, idiosyncratic shocks, not only beyond the origins of these shocks, but also beyond their closest connections as well. Section 1.3.3 tests these factors in detail at the firm level.

In addition, Panels B and C of this table report similar estimates using the 4-quarter growth rate of operating cash flows (Compustat variable OIBDPQ) and change in gross margin

(SALEQ-COGSQ/SALEQ) on the left hand side. The estimates are consistent with the revenue results. In particular, gross margin declines are particularly acute for firms more remote from the origins, potentially indicating that they have lower market power than their suppliers, thus are not as free to adjust prices and consequently are hit with both quantity and price impacts, thereby further contributing to the spillover of shocks.

Finally, note that the previous results are reported as average percentage impacts across all types of shocks. Because different shocks likely have different impact magnitudes, some additional results are in order to provide further clarity on the magnitudes of coefficient estimates. In particular, for each tier of connections (i.e. where $D^n_{i,t} = 1$), I first replace the dummy with the shock's impact on the origin in Regression (1.1), resulting in the following specification:

$$Y_{it,t+4} = a + \sum_{n=0}^{10} b_n Y^n_{0,it} + cX_{i,t} + F_{i,t} + \epsilon_{i,t} \tag{1.2}$$

where $Y^n_{0,it}$ is the 4-quarter revenue growth rate of the firm from which the distance-n shock impacting firm i originates. In this setting, the coefficient estimates can be interpreted as "elasticities": b_n measure the impact of the spillover on subsequent connections $n = 1, \ldots, 4$ in units of the percentage impact on the origin firm. The results are reported in Table 1.5.

[Insert Table 1.5 here]

The results in this table suggest that the spillover effect is not driven only by shocks of large impact. Even after controlling for the impact of the shock on origin firms, the coefficient estimates are still close to 1 at the first three connections and remain statistically significant at distance 4, indicating that shocks that cause a one-unit impact at the origin firm continue to spillover downstream and cause impacts between 40% to 100% of the original impact even to remote, 4th-tier firms.

In summary, these results provide strong evidence that when firms are interconnected in the network of supply chains, shocks previously thought as purely localized are not limited within traditional firm boundaries. Instead, they can transmit along the supply linkages and cause widespread impacts at points remote from origin. I discuss the responses from corporate policies and stock prices in Section 1.5.

1.3.3. Determinants of Spillover Magnitude

Having documented the surprisingly large and far-reaching spillover results, this subsection examines in detail the economic story behind these spillovers. Because existing theories predict that shocks quickly die down after each connection, there must be additional frictions present in the data that give rise to such spillovers. One of them is the departure from perfect competitiveness, which has been discussed in Section 1.3.1 above: if firms located further downstream on average have lower market power than their suppliers, then the revenue impact of the shock could be bigger than predictions made with the assumption of perfect competitiveness. This section empirically examines the existence of these frictions from the data. I first conduct a firm-level test on whether market power is significantly related to the impact of transmitted shock *on each individual firm*. I then conduct a network-level test: I examine the *average* market power at each position of the network, i.e. at each distance from the shocks' origins. In addition, conditional on the occurrence of spillover, input substitutability and inventory levels can further modify its magnitude. I conduct similar firm-level tests for these auxiliary frictions.

Main Channel: Market Power Along the Supply Chain

As discussed in Section 1.3.1 and corroborated by the anecdotal evidence in Figure 1.3, with monopolistic competition, firms' ability to change prices following the spilled over shock beyond perfectly competitive levels can give rise to significant revenue impacts beyond levels predicted by models with perfect competition. At the firm level, this ability is captured by the monopolistic market power of the firm. The spillover effect is particularly acute for competitive firms with less competitive suppliers, which limits their ability to change prices and pass some of the shocks' effects to customers (or final consumers), while enabling their suppliers to pass a larger portion of the impact to them via price increases. I summarize this channel in the following hypothesis:

Hypothesis 2 *Given a shock spilled over from a supplier, the revenue growth rates of firms with lower monopolistic power than their suppliers will be more severely impacted.*

I directly compute the absolute market power of each firm as its size relative to the total size of its most granular industry classification (at the 4-digit SIC level):[9] $MP_{i,t} = \frac{\text{Size}_{i,t-1}}{\sum_k^{N_i} \text{Size}_{k,t-1}}$. I also compute its relative market power to its suppliers as the ratio of its own MP to the average MP of its connected suppliers: $MPR_{i,t} = \frac{MP_{i,t}}{\underset{ji}{E}[MP_{j,t}]}$.

I then test Hypothesis 2 at the firm level using the following regression with MP and MPR measures as interaction variables:

$$Y_{it,t+s} = a + \sum_{n=0}^{4} b_n D_{i,t}^n \cdot MP_{i,t} + \sum_{n=0}^{4} c_n D_{i,t}^n + dMP_{i,t} + \tau X_{i,t-1} + F_{i,t} + \epsilon_{i,t}, \tag{1.3}$$

where D^n are the same shock dummies. To facilitate interpretation, I demean the values of MP and MPR and scale them by the sample standard deviations. The coefficients of interest are b_n and c_n. Here c_n measures the average spillover impact given average values of market power. b_n assesses the validity of Hypothesis 2: it measures the *incremental* effect of one-standard-deviation change in market power (or market power ratio) on revenue growth rate differences between firms with distance-n shocks and firms without shocks.

[Insert Table 1.6 here]

Table 1.6 reports the results. First, the c_n estimates are still significantly negative, albeit smaller in magnitude, in both panels. Therefore, at average levels of market power, shocks can still cause impacts to close and remote firms. Second, the estimates for the interaction term, $D \times MP$, are significantly positive. For example, a one-standard-deviation decrease in the market power of a firm closely connected to the shock origin (n=1) would further decrease the revenue growth by another 25% below average. This is consistent with Hy-

[9]The results are not changed when MP is computed at the 3-digit SIC level and are moderately weaker when computed at the 2-digit SIC level.

pothesis 2 that market power is indeed significantly related to shock spillover at the firm level. Moreover, the estimates for relative market power, reported as $D \times MPR$ in Panel B, are all significantly negative, indicating that firms with lower market power than their suppliers are indeed most exposed to the effect of shock spillovers.

Next I perform the network-level test. Recall that the average spillover effect is significantly negative up to 4 connections away from origin. If this effect arises because of market power, then market power should on average decline as shocks travel downstream from the origin toward final consumers. In other words, on average firms located "more downstream," or close to final consumers, should be more competitive than "upstream firms" located more closely to raw materials. Because the notion of network positions is a relative one, the precise classification of "upstream" versus "downstream" is difficult without individual-good-level data. However, because my data captures the exact origins of each shock, I can proxy the firms' network positions with their distances from the shocks' origins: firms located further away from the shock should on average be more "downstream" and located closer to final consumers.[10] I can therefore conduct a network-level test of the market power channel by comparing the average market power of firms located at each distance from the shock origins.

Starting from the origin of the shock i, I compute the average market power for each distance of connections:

$$\overline{MP_n} = \underset{j}{E}\left[MP_j | j \text{ on path of } i \text{ shock } \& \, Dist(j,i) = n\right], n = 1, \ldots, 10. \tag{1.4}$$

Even though $\overline{MP_n}$ is a relatively crude proxy, if it uniformly declines as n increases, this would provide consistent evidence on the network level that market powers can indeed facilitate shocks spillovers to remote firms. Table 1.7 reports the $\overline{MP_n}$ levels.

[10] In the example from Figure 1.3, distance-2 firms, Dell and HP, indeed disclose that they produce IT products primarily for the consumer market, while the distance-1 firm, Seagate, primarily sells its hard drives to distance-2 firms rather than directly to consumers.

[Insert Table 1.7 here]

This table provide confirmatory evidence that, on average, firms located further down the supply chain from shock origins have lower market power. Because the shocks are distributed across many positions in the network, this measure is relative in nature: origin and distance-1 firms tend to be located more "upstream" in the supply chains and are closer to raw materials, and on average, their market power is significantly higher than distance-3 or -4 firms, which are located closer to the final consumers. The standard deviations are high for all distances, as such, the results cannot statistically establish this pattern for all links. However, they are consistent with the intuition that the more remote, "downstream" firms producing more consumer-oriented goods are on average more competitive than firms producing business-oriented, intermediate goods. Therefore, through the market power channel, spillover to these more remote firms is possible.

Auxiliary Channels: Input Substitutability and Inventories Along the Supply Chain

This subsection links the degree of shock spillover to the degree of substitutability between inputs, which measures the ease of alternative sourcing, and inventories, which serve as a strategic buffer against production interruptions. Consider the same two-tier supply chain (S→C1→C2) mentioned before: first, if individual inputs are perfectly substitutable, then as long as there is more than one supplier at each link, a shock from S would never be able to transmit beyond S, because both C1 and C2 would costlessly switch to alternative sources. At the opposite, suppose that the inputs are perfect complements (i.e. the production functions are Leontief). Here alternative sourcing is impossible, and the shocks would "perfectly" transmit along the chain, with infinite marginal impact to both C1 and C2. Real-world production probably falls somewhere between the two extremes, while Boehm et al. (2014) find evidence that the relation is close to Leontief between imported and domestic inputs. Therefore, while not enough to generate spillover on their own, variations in input substitutability can further modify the magnitude of spillovers, conditional on their occurrence: if the first-tier (S→C1) connection has higher substitutability than higher-tier (C1→C2) connections, the effect on C2 could be bigger than predictions made using, say, Cobb-Douglas production functions.

Similarly, when hit with a supply shock, firms with more input inventories already have more redundant inputs at hand, so they are less subject to the impact of the shock. Firms with infinite inventories, similar to those with perfectly substitutable inputs, are never subject to any spilled over shocks. Firms with zero inventories are the most sensitive to the impact of any shock. Inventories are costly to hold, so most real-world firms hold some minimum "safety level" computed using common operations models. If the average level of inventories differ at different distances from the origin of supply shocks, then *ceteris paribus*, this would result in different levels of shock impact on average. I summarize these channels in the following hypothesis:

Hypothesis 3 *Given a shock spilled over from a supplier, the revenue growth rates of firms whose suppliers are less substitutable, and firms with lower inventory levels, will be more severely impacted.*

A firm-level proxy for input substitutability is needed for the empirical tests. For a smaller sample of firms, my data capture exact value of the supply relationships as a fraction of the customer's inputs, i.e. $\gamma_{ji,t} = \frac{V_{ji,t}}{COGS_{i,t}}$. If supplier j constitutes an important part of customer i's production, as indicated by a high γ_{ji}, then j is likely harder for i to substitute. In this setting, the γ_{ji}s measures the "switching cost" of finding alternative suppliers in the event of a disruptive supply shock. As such, they directly proxy for the substitutability of each supplier. For each customer i, I compute its average supplier substitutability as $\overline{\gamma_{i,t}} = \underset{ji}{E}\left[\gamma_{ji,t}\right]$. To ensure robustness, I use several industry-based alternative measures of substitutability and report their results in Appendix A.1.4. Firm-level inventories are defined as the ratio of total inventory to total assets: $INVR_{i,t} = \frac{INVT_{i,t-1}}{AT_{i,t-1}}$.

I then test Hypothesis 3 with the following regressions:

$$Y_{it,t+4} = a + \sum_{n=0}^{4} b_n D_{i,t}^n \cdot \overline{\gamma_{i,t}} + \sum_{n=0}^{4} c_n D_{i,t}^n + d\overline{\gamma_{i,t}} + \tau X_{i,t-1} + F_{i,t} + \epsilon_{i,t}, \tag{1.5a}$$

$$Y_{it,t+4} = a + \sum_{n=0}^{4} b_n D_{i,t}^n \cdot INVL_{i,t} + \sum_{n=0}^{4} c_n D_{i,t}^n + dINVL_{i,t} + \tau X_{i,t-1} + F_{i,t} + \epsilon_{i,t}. \tag{1.5b}$$

where I demean the values of $\overline{\gamma}$ and $INVL$ and scale them by the sample standard deviations. The interpretation of these regressions is the same as Regression (1.3): b_n assesses the *incremental* effect of one-standard-deviation change in supplier substitutability (or inventories) on revenue growth rate differences between firms with distance-n shocks and firms without them. The coefficient estimates are reported in Table 1.8.

[Insert Table 1.8 Here]

The roles of inventories and supplier substitutability are evident from this table. The estimates for both $D \times INVR$ and $D \times \overline{\gamma}$ are significantly positive and economically large. They can be interpreted as follows: when a firm experiences a shock, possibly from a remote source along its supply chain, if its inventory level, or supplier's substitutability, is one standard deviation higher/lower, then the spillover effect would be mitigated/exacerbated by 11% and 37% of the average revenue growth rates, respectively. Therefore, conditional on a shock occurring from somewhere in the supply chain, both inventories and supplier substitutability can indeed further affect its impact on the firm.

A word of caution is that, when put in the same regression, only MP (or MPR) remains statistically significant. This is because market power leads to impacts on both prices and quantities, while inventories and substitutability likely affect quantities only. Therefore, these two channels are likely auxiliary to the market power channel.

1.4. Key Robustness Tests

The surprisingly pronounced spillover effects documented in the previous two sections demand careful checks and examinations in order to firmly establish their internal and external validity. The unique, granular nature of the data facilitates a series of direct robustness tests. In this section I discuss several tests that are most helpful in solidifying the paper's results. A host of additional robustness checks can be found in Appendix A.1.4.

1.4.1. Are the Shocks Well-Identified?

Recall that the shocks captured by my data have to meet two hurdles to be considered well-identified idiosyncratic shocks: they have to have both idiosyncratic causes and localized, firm-specific effects. The construction and summary statistics presented in Section 1.2.1 have provided evidence on the first condition. This subsection specifically discusses tests needed to confirm the second condition.

Some Large Shocks Might Have Systematic Effects

The first concern is perhaps the most salient: some shocks, such as large natural disasters, although idiosyncratic in nature, might have wide-ranging impacts on not only the origin firms that report these shocks, but also their connected customers, either directly (i.e. the disaster hits both the disclosing firm and their customers) or indirectly (i.e. the disaster creates a region-wide demand shock that feeds back to the revenue of both origin firms and their customers). In both cases, the shock's spillover effect through linkages is contaminated by the "systematic effect" of the shocks and the coefficient estimates in my empirical framework would be amplified. Therefore, if on average my sample is populated by these large shocks, then the empirical results cannot be ascribed solely to shock spillover through supply chain linkages.

Fortunately, because LDA classifies shocks into distinct categories, I can perform three tests to rule out this case. First, I replicate the spillover results from Regression (1.1) while *deleting* from the sample one shock type at a time. I then redo the tests *using* only one category at a time. Finally, I keep only one category of shocks that are unequivocally idiosyncratic in both cause and effect: *fires*, which are extracted from the manmade disaster category of my sample by a simple keyword search for fire-related words. This results in a collection of 174 shocks, and I redo the spillover analysis using only this reduced sample. The results for these tests are reported in Panels A to C of Table 1.9.

[Insert Table 1.9 here]

In all three cases the results are not significantly changed, which confirms that on average, my sample is not contaminated by shocks with potentially systematic effects. Their initial impacts are localized to the origin firm only, thus their subsequent impact on connected customers are indeed due to the spillover induced by the supply chain linkages.

Prior Growth Trends

Next, for the disclosure-based shocks to be well-identified idiosyncratic shocks, before their onset, a firm's revenue growth should be similar to those that are never hit with any shocks. This implies that there must *not* be a statistically significant difference in prior-period growth rates between these two groups. To test this prior trends assumption, I fit Regression (1.1), replacing the left hand variable with lagged one quarter $(t-1,t)$ to 8 quarters $(t-8,t)$ revenue growth rates, for both the origin firms $(n=0)$ and their subsequent connections $(n=1,\ldots 4)$. I tabulate these results in Table 1.10.

[Insert Table 1.10 here]

This table clearly demonstrates that none of the estimates for $(t-1,t)$ to $(t-8,t)$ are statistically significant. This indicates no significant difference in prior-to-shock revenue growth rates between treated (hit with shock) and control firms in the previous two years. As such, the significant difference in growth rates shown in Table 1.4 are likely due to the spillover effect through supply chain linkages rather than any existing trends in the data.

Strategic Disclosure of Shocks by Firms

A key feature of my shock data is that they are extracted from voluntary firm disclosures, such as 8-Ks and press releases. Therefore, one might worry whether firms are completely truthful in their disclosures, or if there are systematic differences in the reporting standards of these events.

This concern is legitimate and important, because strategic disclosure of these events can introduce amplifying biases to the estimates from the spillover regressions, for three reasons: First, firms might disclose a shock only if its impact is too large to strategically hide from public view. In this case, the shocks that I capture would be overrepresented by large idiosyncratic shocks, thereby artificially inflating the true spillover effect from the interconnections.

A closely related concern is reverse causality: A manager facing low future revenue growth rates might attempt to "explain away" this bad performance by pointing to "idiosyncratic shocks to the supply chain" that are out of their control. Zhou (2014) documents the usage of such "external blame" language in earnings conference call transcripts. If such strategic blame is prevalent in the shock disclosures, then the true spillover effects measured in Regression (1.1) could also be contaminated upwards.

I first address the concern of unbalanced reporting of only large shocks, via an exogenous change in the disclosure reporting standards. In particular, an extension to the Sarbanes-Oxley Act (SOX Act) mandates additional "real time" disclosures of information on material changes in firms' *operations*, which include supply chain-related activities, beyond the original SOX Act requirements. The SEC began to strictly enforce the provision on August 23, 2004.[11] This enforcement thus serves as an exogenous policy shock after which firms would presumably disclose more supply disruptions, which is confirmed by Schmidt and Raman (2013). I therefore cut the sample into two halves corresponding to before and after

[11] https://www.sec.gov/rules/final/33-8400.htm

this enforcement date, and replicate Regression (1.1) on each subsample. If unbalanced reporting is a problem, then the results would be weaker in the post-2004 subsample. Panel A of Table 1.12 presents the results. Results before and after 2004 are very similar in both economic magnitudes and statistical significance, indicating that firms are not intentionally disclosing only major shocks.[12]

[Insert Table 1.12 here]

Next, the reverse causality concern is addressed with subsample analysis. Recall that my sample shocks are extracted from two types of disclosures: 71.03% are shocks to the disclosing firm itself and 28.97% are disclosed shocks from external suppliers. Because strategic blaming of external suppliers is not an issue for the former sample, I replicate Regression (1.1) using only this restricted sample of shocks and report the coefficients in Panel B of Table 1.12. Here again, the coefficients are not significantly different between the restricted and unrestricted samples, suggesting that reserve causality related to the strategic blaming of external parties in disclosures is not significantly present in my shock data.

Finally, recall Table 1.1 before: it also shows that the firms disclosing shocks are not statistically different in size, book-to-market ratios, and other dimensions, from those without disclosures. This indicates that the disclosure decision of a shock is not related to any particular firm characteristics, alleviating potential concerns of some firms strategically reporting these shocks more than others.

Overall, the tests in this subsection further confirm that my sample of LDA-classified firm disclosures capture well-identified idiosyncratic shocks at the firm level. The versatility of my methodology enables its use to identify shocks in more general contexts as well (e.g. shocks outside supply chains such as financing shocks), as a robust identification tool for empirical finance researchers.

[12] Inserting the post-enforcement period as a dummy variable produces similar results.

1.4.2. Are Shocks Correctly Mapped to the Network?

Having established that the shock data are well-identified and idiosyncratic, one still needs to make sure that these shocks are correctly mapped to the network, and that the spillover results are indeed caused by these mapped shocks and not by 1) spurious statistical factors and 2) the mathematical operation of taking averages on network links that are unevenly distributed at different positions. The first concern is easily addressed through a falsification test discussed below. I defer the technical details of the second test to Appendix A.1.4 for interested readers.

Because my data record the date of occurrence for each shock, I conduct the following falsification test to rule out spurious relations. Specifically, for each shock date t and firm i, I first reset all $D^0_{i,t} = 0$. I then randomly assign falsified "shocks" to firms on each occurrence date, i.e. $\exists\, \hat{D}^0_{i,t} = 1$ for some random i. I then trace out the customers of these firms, recreate similar tiered dummies $\hat{D}^n_{i,t}$. $n = 1, \ldots, 4$, and repeat Regression (1.1). Table 1.11 reports the coefficient estimates.

[Insert Table 1.11 here]

This table demonstrates that, while using real data, the spillover effect is significantly negative for the first 5 distances; none of the estimates are statistically significant using randomly assigned shock data. Therefore one can be more confident that the spillover impacts are indeed caused by the mapped shocks to the network linkages, and not by spurious relations.

1.4.3. External Validity: Private Firms

A valid concern about the spillover results from the previous sections is their generalizability: crucially, my data contain only publicly-traded firms. This leads to two potential concerns, first, would network data be general and complete enough even with the lack of private firms? Second, would omitting private firms lead to amplifying biases on coefficient estimates? If both can be ruled out, then the spillover result, even though derived using public-firm data, can be readily generalized to all firms.

Recall that a subset of my network data contains valued relationships. This enables me to empirically examine the first concern. The intuition is as follows: first, the make-use table between sectors (defined by NAICS codes) constructed by the BEA contains data from both public and private firms with more than 50 employees. Second, if I aggregate the relationship values flowing from all firms in one sector to all firms in another, I can reproduce a similar table and compare the sectoral-level input shares $V/COGS$ with the shares in the BEA table. Therefore, if the shares are comparable, then adding or removing private firms probably would not fundamentally change the shape of the overall network.

I aggregate the smaller sample of 7,054 domestic firms with valued supplier relationship data into 41 industry groups according to BEA's definition.[13] I then compute the within- and between-group trade shares by summing up the values at the end of 2013 for all firms within and between each group, then scale by the sum of COGS. This results in 1,681 sectoral shares, which I then correlate with the same shares computed from the 2013 use table from the BEA.[14] Appendix A.1.2 details the construction steps. The Spearman's rank correlation between these shares is 0.815. This indicates that, at least on the sectoral level, my public-firm data produce a network of similar shape to the combined data of both public and private firms.

Second, because my shocks are all originated from publicly traded firms, to the extent that private firms serve as alternative suppliers to other public firms along the shocks' paths, this would produce an attenuation bias to the estimates, making it more difficult to find a significant result. One would also raise other concerns related to external validity, e.g. the generalizability of results to positive shocks. I address these additional concerns in

[13]The actual BEA table has 71 industries. I remove industries that are not present in my dataset such as financial services, resulting in 41 groups.

[14]Available at
http://www.bea.gov/industry/xls/io-annual/IOUse_Before_Redefinitions_PRO_1997-2013_Summary.xlsx

Appendix A.1.4.

1.5. Financial Response from Firm and Market to Spillovers

The results in the previous sections document a pronounced and robust spillover effect in the network of supply chains: localized, firm-specific shocks can cause substantial impact to the revenue of firms located up to 4 connections away. Given both the broad reach and the large magnitude, one would intuitively expect significant responses, from both corporate financial policies and firm valuations, to these significant spillovers.

However, existing theories and empirical research have not explored these responses. Specifically, first, do firms even recognize this spillover and adjust anything in response? Second, what policies would firms adjust? Third, how much would they adjust? Do they simply make up shortfalls after these spillovers, or would they respond more pro-actively and aggressively? Finally, is there any heterogeneity in the response to close vs. remote spillovers?

By examining firms' post-shock discussions from their 10-K/Qs, this section first uncovers a series of policies that the firms *say* they would adjust, then empirically tests whether firms actually *do* adjust these policies and the intensity of such adjustments. I then examine the adjustment in firm values by documenting the stock price response to these shocks, and uncover heterogeneities in responses from shocks spilled over from close vs. remote sources.

1.5.1. Hypothesis Development

Because previous research has not provided empirical evidence or theoretical insights on corporate responses to supply chain shocks, one needs to be careful in the analysis to avoid data snooping bias. Fortunately, the same firms that disclose these shocks often discuss their response to these shocks in subsequent disclosures such as 10-K/Qs. Therefore, from these disclosures, I can directly extract the list of variables that are most often talked about.

I proceed with a simple three-step keyword search: First, similar to Jegadeesh and Wu (2013), for each firm that discloses a shock, I obtain their 10-K/Qs filed in the next four quarters, and isolate the Management Discussion & Analysis (MD&A) section. Second, from these discussions, I use the same keyword filters discussed in Section 1.2.1 to look for sentences related to these shocks. Third, from each keyword within the sentences, I tabulate the frequency of the words closest to the keyword and manually isolate the most frequent occurrences related to corporate policies. This results in three groups of top keywords related to:

1. Strategic buffers and reserves: *cash, working capital, inventory, buffer, excess*, etc
2. Production adjustments and alternative sourcing: *redesign, accommodate, alternative supplies, redundant sourcing*, etc
3. Financing motives: *strong balance sheet, funding source*, etc

These words directly give a starting point for the empirical analysis. First, they relate to buildups in working capital such as cash and inventories, possibly beyond the levels before the shocks' impact. Second, they relate to adjusting production or redesign some aspects of products to accommodate a wider selection of alternative suppliers, which suggest changes in capital expenditures and possibly R&D investments into these adjustment costs and technologies. Third, firms seem to be concerned about how to finance these buildups, indicating that they might tap into external funding sources such as debt or equity issuances. I summarize these insights in the following hypothesis:

Hypothesis 4 *Following a spilled over shock, firms enter an active buildup phase in their working capital and longer-term investments, possibly beyond simple recovery levels. Such activities require financing beyond internal funds.*

I test this hypothesis using the following regressions:

$$CF_{it,t+k} = a + b_n D^n_{i,t} + cX_{i,t-1} + dF_{i,t} + \epsilon_{i,t},\ (n = 0, 1, 2, k = -1, \ldots, 8), \tag{1.6}$$

where D^n is the same shock dummy from previous specifications, except that I now group all firms beyond the closest connection into the "remote connections" group and set $D^{n=2} = 1$ $\forall n \geq 2$. CF is the change in corporate investment and financing policies including:

- Working capital: (1) Cash, (2) Inventories
- Capital expenditures: (3) CAPEX, (4) R&D expense
- Financing "pecking order": (5) Retained earnings, (6-7) net equity and debt issuances, and (8) trade credits (account payables)

All CF variables are changes scaled by lagged total assets. The exact construction for each variable can be found in Appendix A.1.1. Here the b_n coefficient measures the incremental change in corporate policies for firms hit with shocks compared to those that are not.

1.5.2. Proactive Cash and Capital Buildups

I first fit Regression (1.6) with the first four variables related to working capital and capital expenditures on the left hand side. To rule out prior trends, I first use changes from the $(t-1, t)$ period. I then use changes from $(t, t+1)$, $(t, t+4)$, and $(t, t+8)$ quarters. Table 1.13 reports the results.

[Insert Table 1.13 here]

Three observations are immediate from this table. First, contemporaneously, the shocks have negative impacts on the working capital of both origin firms and their customers: in the second row of this table both inventories and cash holding changes are from 2.0% to 8.3% standard deviation lower than comparable firms. In the quarters after the shock, however, firms hit with shocks (either directly or through spillovers) in turn accumulate significantly more cash and inventories than those not hit with shocks: rows 3 and 4 of the table indicate that cash and inventory growth are 4.7% to 14.8% standard deviation higher. This result strongly suggests a significant buildup in working capital after the shocks. Perhaps in a similar mechanism as Dessaint and Matray (2015), firms hit with supply shocks become more salient, causing managers to engage in more precautionary behaviors.

Second, the buildup behavior is also evident in longer-term capital expenditures: CAPEX growth ranges between 2.6% and 5.8% standard deviation higher during the next 8 quarters, while R&D expense is mildly higher but not as statistically significant. This is consistent with firms' own disclosures that they pro-actively seek to accommodate alternative supply sources after the shocks, by reconfiguring their production systems or redesigning some aspects of their products. These efforts would, and indeed have, translated into higher capital expenditures.

Third, interestingly, the buildup is much more muted for firms experiencing shocks spilled over from more remote sources (distance n¿1). For these firms, only CAPEX is significantly higher than comparable firms, and the buildup in working capital from closer connections is absent. This suggests the existence of information acquisition frictions within the deeper structures of supply chains: as the supply chains become longer, it would be exponentially more costly for firms to accurately monitor and trace the source of shocks from the more remote tiers of connections. As such, their response to remote shocks is more muted and slower. Anecdotally, this "limited visibility" has been extensively discussed in firms' disclosures, and the empirical results from this section can motivate a class of information-theoretic models where the network structure of inter-firm relations can be explicitly accounted for, and thus generate heterogeneous responses to shocks from different network positions.

1.5.3. Changes in Capital Structure and Financial Leverage

The capital buildup documented in the previous subsection is substantial in magnitude. Therefore, both the researcher, and indeed the firms themselves, would be concerned whether the firms have enough internal resources to finance these endeavors, or need to tap into external sources of financing. I therefore use the last four variables in the list from Section 1.5.1 on the left hand side of Regression (1.6). These are related to the "pecking-order" of financing, and the coefficient b_n measures the incremental change of these financing policies after the shock spillover, compared to firms not hit with any shocks.

[Insert Table 1.14 here]

I report the results in Table 1.14 above. Similar to the previous table, three observations are evident from Table 1.14: first, the buildup documented in the previous subsection does not seem to be financed by internal funds. This is perhaps expected because the shocks captured by the disclosure data are mostly negative, leading to lower operating cash flow growth (from Table 1.4) and, *ceteris paribus*, lower retained earnings, which is confirmed in this table as changes in retained earnings range between 2.4% to 6.5% standard deviation lower. Similarly, because these shocks are shocks from suppliers, it is not surprising that trade credits are also slightly lower.

This leaves external sources of funding such as debt or equity issuances. The last three columns of the top panel show little change in equity issuances post-shock. However, debt issuance is significantly higher for both origin firms and their close, and remote, connections, ranging from 1.4% to 4.7% standard deviation higher. This translates into market leverage that is between 2.2% and 9.4% standard deviation higher. These results suggest an interesting trade-off: firms are taking on higher debt levels to build up their working capital reserves, which is itself a costly endeavor. As such, shock spillovers through the supply chains lead to firms trading off higher default risk (as a consequence of becoming more levered) with lower operating risk (mitigated effect of spillovers). This again leaves room for future theory development in financing policies.

Third, similar to the previous table, the increase in debt issuance and leverage is also more subdued for firms facing more remote shocks. This again suggests the existence of information frictions that are severe and costly enough to prevent firms from engaging in potentially costly increases in financial leverage without clear visibility into the exact sources of spillovers.

1.5.4. Changes in Firm Valuation

This subsection examines the effect of shock spillovers on firm value. I form my hypothesis as follows: First assume that the market is not aware that shocks have hit the firm before they are disclosed (my tests also check for pre-announcement leaks). Then, when a firm discloses these shocks, an efficient market should immediately deduce the magnitude of the effects, and discount the stock prices accordingly to reflect the updated firm valuation. Because my shock sample consists of disclosures of both firms' own shocks and shocks from their tier-1 suppliers, I form the first equal-weighted stock portfolio consisting of all origin firms and their immediate customers i.e. the S and C1 firms in the S$\rightarrow$C1$\rightarrow$C2$\rightarrow \ldots$ chain.

The effect on the remote connections, i.e. C2, C3...'s price, however, is uncertain. In the extreme case, the market might not even know that C1 and C2 are linked. A more likely scenario is that the market is aware of the linkage, but takes some time to fully ascertain the effect as C2 is further away from the shock source. This can either be due to investor inattention as suggested by Cohen and Frazzini (2008), or due to information processing constraints that prevent investors from processing complex network structures in a timely fashion. Therefore, the abnormal market return for remote connections, if any, might persist for a longer period than that of the origins and their immediate connections. I summarize this intuition in the following hypothesis:

Hypothesis 5 *Firms that experience a supply disruption should experience negative abnormal returns around the event's reporting date. If customers linked to the firm also experience negative abnormal returns, then these returns should persist for a longer period after the reporting date.*

For a window of [-10, 40] trading days around the report date t, I define the report period cumulative abnormal return (CAR) and abnormal turnover (AT) as follows:

$$CAR_{i,t+s} = \prod_{k=-10}^{s} Ret_{i,k} - \prod_{k=-10}^{s} Ret_{vw,k},$$
$$AT_{i,s} = \frac{60Vol_s}{\sum_{k=40}^{100} Vol_{s-k}} - 1,\ s \in [-10, 40], \tag{1.7}$$

where $Ret_{i,t}$ and $Ret_{vw,t}$ are gross returns on stock i and on the CRSP value-weighted index on date t. Vol_t is the trading turnover on date t. CAR measures market-adjusted returns while AT measures the extra trading volume as a fraction of the previous 60 trading days from $t-100$ to $t-40$. Figure 1.4 plots the average CAR and AT for the three equal-weighted portfolios consisting of (1: S+C1), (2: select C1), and (3: C2,C3...) firms, respectively.

[Insert Figure 1.4 here]

Three observations are evident from this figure. First, the reported events captured by my dataset are not leaked in advance. Both abnormal returns and abnormal turnover prior to the events' disclosure dates are not significantly different from zero. The events therefore are likely reported on a timely basis consistent with SEC-mandated disclosure standards, and do not represent stale news either.

Second, the solid black line of this figure (left axis) plots the cumulative abnormal returns (CAR) for the equal-weighted portfolio consisting of all origin firms and their tier-1 customers. Here, consistent with the hypothesis, the market reacts promptly to supply shocks on the origin firms and their immediate customers: For the three days [t,t+2] after the event, the cumulative abnormal return is -3.98%. The return does not revert in the following days, indicating that the market is indeed cognizant of the first-tier spillover effect of the shock on real outcomes such as revenue.

Third, the dotted black line plots the CAR for the equal-weighted portfolio of all remote connections of the origin firms beyond the tier-1 connections. The market does not seem to immediately price in these remote, spilled-over shocks, and the CAR declines more slowly, and persistently drifts downwards for up to 40 trading days. This indicates that the market is slower to realize the impact of shock spillovers from more remote sources. In addition, the gray line plots the CAR for the equal-weighted portfolio of all tier-1 customers that *did not directly report* any supplier shocks. That is, for these firms, the impact has to be inferred from someone else's disclosure, either the supplier's or from another customer that is also connected to the supplier. Here the market reaction is also slower than the case of direct disclosures, consistent with the information processing constraint channel: inferring these indirect network linkages takes time, and the market is thus slower in fully adjusting the stock prices for firms located more remotely from the origins.

These results suggest the existence of possible profitable arbitrage opportunities for investors who are more adept at analyzing the complex structure of supply chains, and also establishes the linkage of the spillover effects to network-related systematic risks and expected stock returns. I explore these asset pricing related issues in a related paper.

1.6. Concluding Remarks

This paper is the first in finance to empirically quantify how firms' localized, idiosyncratic shocks spill over along the supply chain interconnections and affect both close and remote firms along the chain. This is achieved via precisely mapping 1) a hand-built database of over 8,000 firm-level idiosyncratic supply shocks of diverse types, extracted from the texts of over 5 million firm disclosures, to 2) a hand-built network of over 1 million supply chain connections between publicly traded firms globally.

The results suggest that when firms are interconnected in the complex web of supply chains, localized shocks are not that local: they cause substantial impact to firms even up to 4 connections away from the origin, through frictions such as the uneven distribution of market power along the supply chains. Facing such significant risk propagation, managers respond with significant buildups in capital financed by debt issuances, leading to higher financial leverages. The stock market reacts to shock spillovers from distant connections with slower speeds: post-shock abnormal returns are persistently negative for up to 40 days. The empirical results in this paper provides the economic foundation for future theory developments on production networks and asset pricing. I explore some of these issues in related works.

Figure 1.1: Visualization of Sample Supply Chain Network Shape, 2002 and 2015

This figure presents a visualization of a portion of the network: 400 select US firms from technology-related industries (Fama-French industry codes 35 to 37) in the years of 2002 and 2015. The network data is constructed from a combination of firm disclosures and proprietary data sources described in detail in Section 1.2.2.

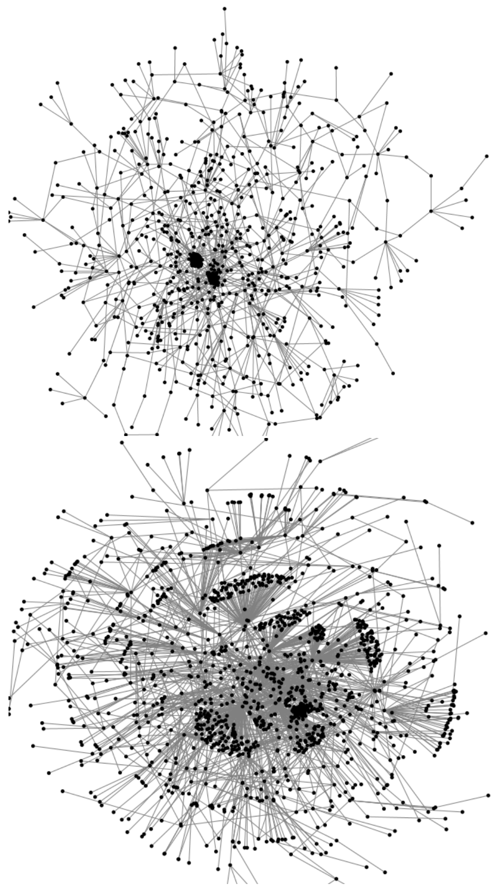

Figure 1.2: Distribution of Shocks Across Time and Types

The top panel of this figure plots the number of total shocks captured by the disclosure data over time. The bottom panel of this figure plots the number of each type of idiosyncratic shocks, as classified by the LDA, over time from 1994 to 2015. The shock classification procedure is described in Section 1.2.1 of the text.

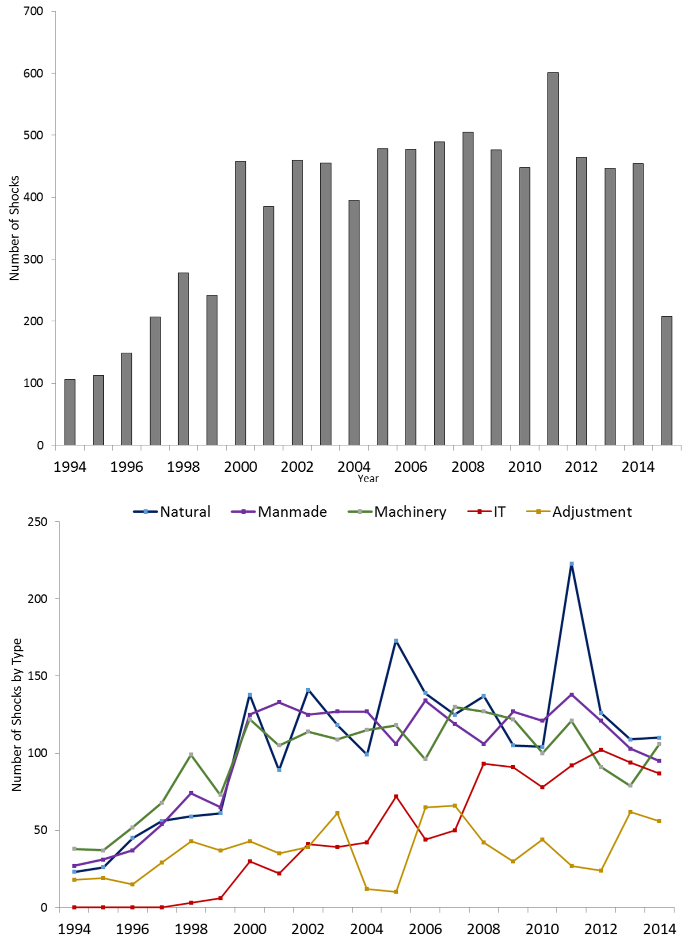

Figure 1.3: Example Timeline of A Shock Spillover

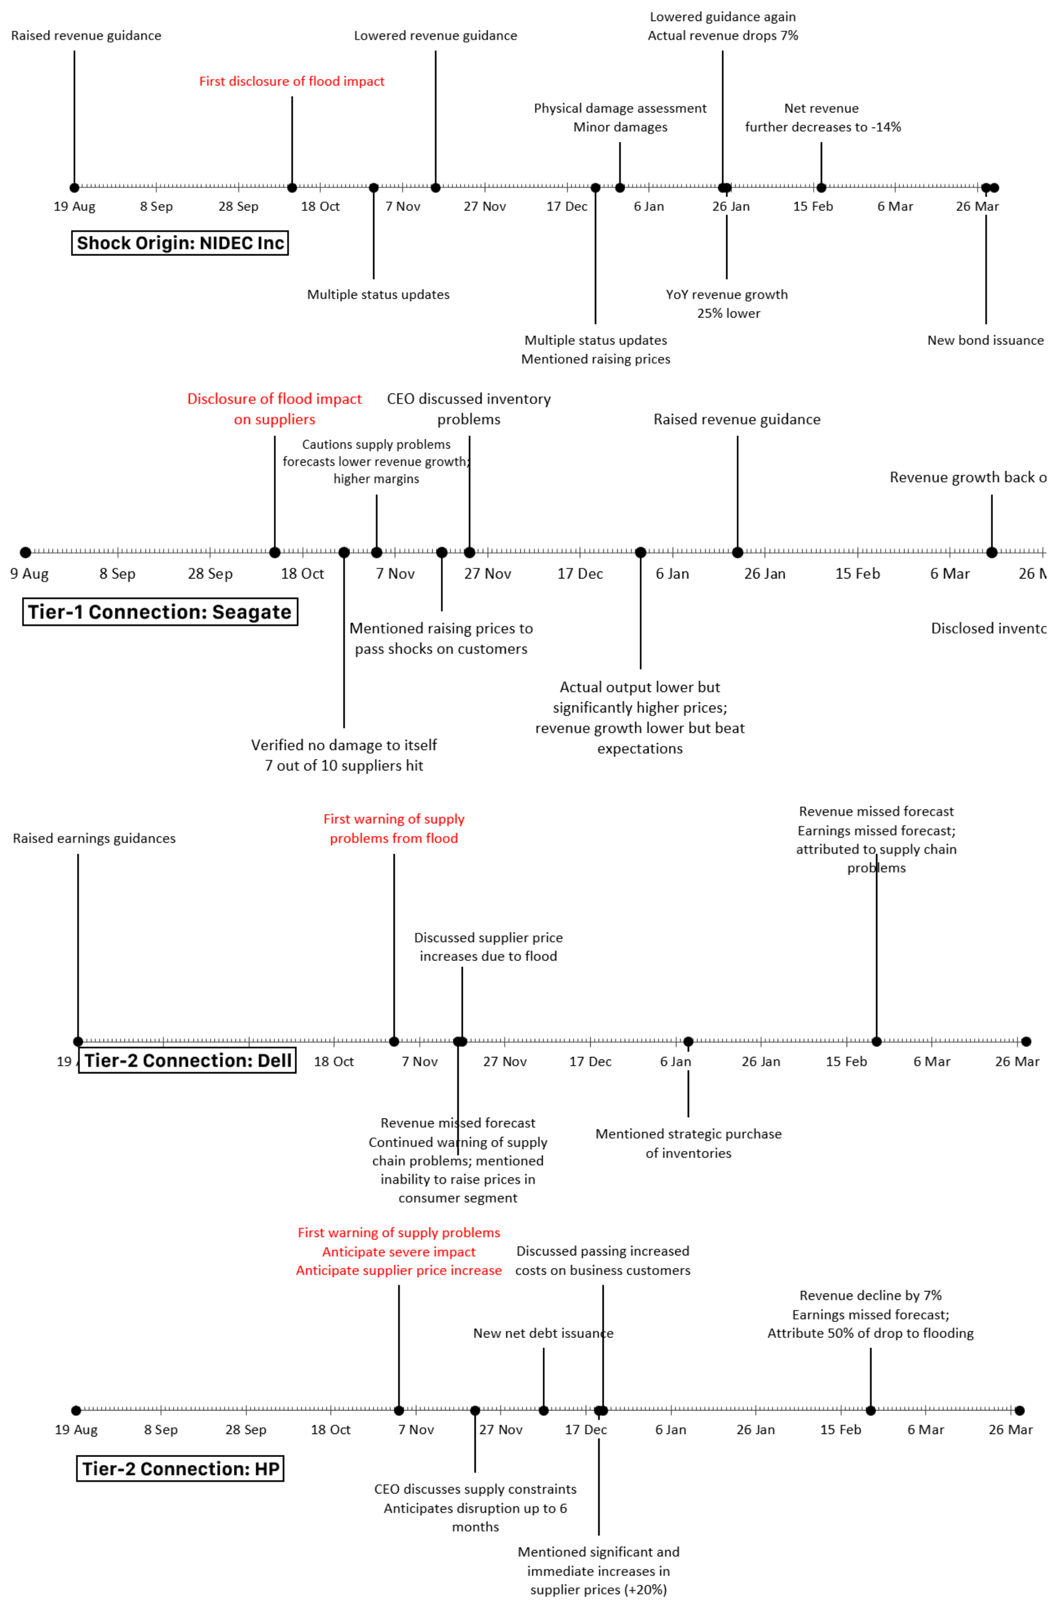

Figure 1.4: Stock Price Reactions Following Close vs. Remote Shocks

The solid black line of this figure (left axis) plots the cumulative abnormal returns (CAR) for the equal-weighted portfolio consisting of all origin firms and their tier-1 customers, around the date when a supply shock is first disclosed. The gray line plots the CAR for the equal-weighted portfolio of all tier-1 customers that did not directly report any supplier shocks, The dotted black line plots the CAR for the equal-weighted portfolio of all remote connections of the origin firms beyond the tier-1 connections. The light gray bars (right axis) plot the daily abnormal turnovers (AT) for the equal-weighted portfolio consisting of all origin firms and their tier-1 customers, and the dark gray bars plot the average AT for all remote connections of the origin firms beyond the tier-1 connections. The CAR and AT measures are computed according to Equation (1.7) of the text from 10 trading days prior to 40 days after the shock's disclosure date. The negative abnormal turnovers are truncated at -1% to save display space.

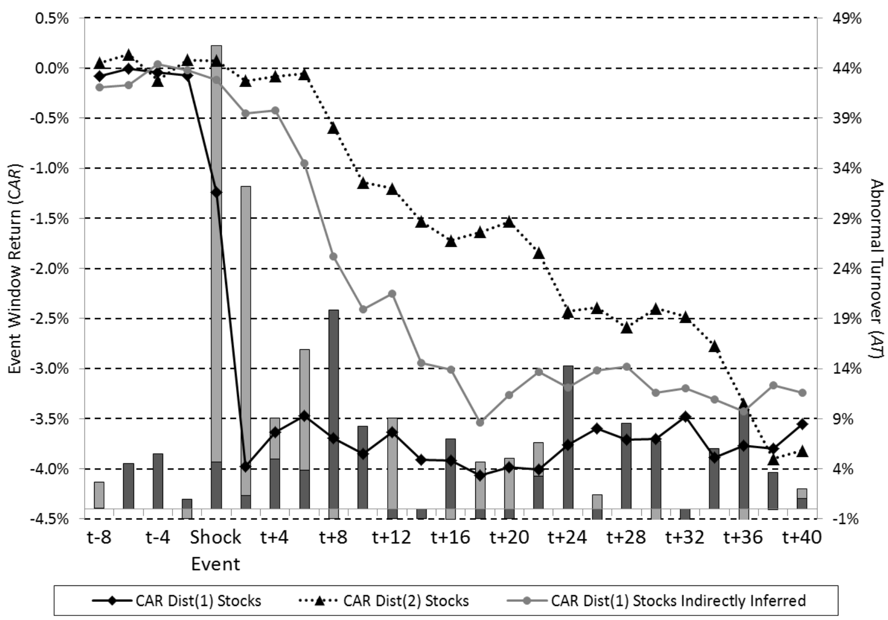

Table 1.1: Summary Statistics of Idiosyncratic Shock Data

This table presents summary statistics of supply chain shock events extracted from firm disclosures from 1994 to 2015. The top panel reports sample sizes and number of firms. The middle panel reports the breakdown of captured shocks by major type as classified by the LDA algorithm, discussed in Section 1.2.1 of the text and Appendix A.1.2. The bottom panel compares characteristics such as firm size, book-to-market P/E ratio, return on assets, leverage ratio, and inventory level, between disclosing vs. non-disclosing firms. The numbers in brackets are *t*-statistics for the quarterly average difference in quarterly levels of these measures between disclosing firms and the overall sample. The definition and construction of all variables can be found in Appendix A.1.1.

Panel A: Overall Shock Statistics			
Total Number	11191		
% of firms reporting at least one shock	20.06%		
Avg no. of shocks per firm	5.104		
No. of shocks with supplier identified	8295		
No. of shocks matched to network data	8160		
Panel B: Distribution of Shock Types			
Types of Identified Shocks	# of Events	Percent sample	
Natural disasters	2256	27.20%	
Manmade disasters	2145	25.86%	
Production glitches	2076	25.03%	
IT issues	1032	12.44%	
Adjustment failures	786	9.48%	
Total	8295	100.00%	
Panel C: Summary Statistics of Disclosing Firms vs. Full Sample			
Average	Reporting	Diff	t-Stat
Size	2.201	0.156	(1.94)
BM	0.687	-0.028	(-0.27)
PE	13.902	0.677	(0.68)
ROA	0.087	-0.023	(-0.56)
Leverage	0.411	0.039	(0.13)
Total Inventory	0.148	0.013	(1.20)
No. Obs (quarters)		82	

Table 1.2: Top 5 Keywords for LDA-Identified Shock Topics

This table reports the top 5 words for each topic identified by the LDA procedure. Each column in this table represents a topic $k = 1, \ldots, 20$.

Group 1: Systematic Types

Topic 1	Topic 2	Topic 3	Topic 4	Topic 5	Topic 6
global	uncertainty	economy	consumer	sector	retail
systematic	global	condition	economic	industry	distributor
markets	risk	recession	demand	competitive	sales
widespread	terrorism	condition	capacity	cost	seller
countries	property	growth	consumption	price	third-party

Group 2: Middle Types

Topic 1	Topic 2	Topic 3	Topic 4	Topic 5	Topic 6
worker	union	government	research	transportation	quality
labor	strike	legal	intellectual	channel	design
strike	organization	regulation	property	logistical	warranty
stoppage	wage	licence	dispute	development	flaw
employee	relation	regional	restriction	oursourcing	recall

Group 3: Idiosyncratic Types

Topic 1	Topic 2	Topic 3	Topic 4	Topic 5	Topic 6	Topic 7	Topic 8
disaster	flood	fire	hurricane	machinery	breakdown	IT	failure
destruction	water	damage	weather	equipment	equipment	breach	install
earthquake	recovery	accident	tornado	production	assembly	information	equipment
damage	damage	power	storm	suspend	factory	sensitive	manufacturer
catastrophe	disaster	electricity	sustain	shutdown	outage	intrusion	maintainance

Table 1.3: Summary Statistics of Sample Firms and Network Connections

A: *Number of Firms*	Total	Domestic	Foreign			
Total in Sample	10930	7089	3841			
B: *Overall Sample Statistics*	Mean	Median	75p	25p		
Revenue Growth	0.107	0.071	0.217	-0.044		
Size	2.045	2.577	12.212	0.580		
BM	0.715	0.552	0.909	0.314		
PE	13.225	14.971	23.217	-1.823		
ROA	0.110	0.106	0.139	-0.004		
Leverage	0.372	0.138	0.457	0.009		
No. Obs		335337				
C: Treated vs. Untreated	Distance from Shock					
	0 (Origin)	1	2	3	4	¿4 or Never
Size	2.201	2.218	1.954	1.819	1.911	2.073
BM	0.687	0.682	0.802	0.779	0.570	0.703
PE	13.902	12.871	14.043	13.198	12.984	12.981
ROA	0.087	0.108	0.130	0.105	0.109	0.112
Leverage	0.411	0.371	0.394	0.335	0.404	0.368
Inventory	0.148	0.139	0.101	0.082	0.148	0.153
No. Obs	8160	36477	40469	44290	45491	116044

D: Average Link Statistics			
Total Sample		Quantified Sample	
Total # of Links	1007998	# of Quantified Suppliers	24.69
Total # of Quantified Links	473759	# of Quantified Customers	18.65
Total Links Per Firm	92.22	Mean Total Supplier share	0.34
Quantified Links Per Firm	43.34		

Table 1.4: Spillover of Idiosyncratic Shocks in the Network: Average Results

This table reports the coefficient estimates of b_n, $n = 0, \ldots, 4$ from Regression (1.1) of the text. b_n measures the average difference between firms hit with a shock spilled over from a distance of n connections, and firms never hit with any shocks. The dependent variables in Panels A to C are growth rates in revenue, operating income, and gross margin, respectively. D^n is a dummy variable that equals to 1 if one of firm i's suppliers from a distance of n connections experiences an idiosyncratic shock captured by the disclosure data. All control variables are defined in Appendix A.1.1. All standard errors are clustered at the firm level. All regressions include industry×year, fiscal quarter, and state/country fixed effects, and are in quarterly frequency from 1994 to 2015.

Panel A: Four-Quarter Revenue Growth Rates					
	Distance from Shock Origin (in # of Connections)				
	Origin	n=1	n=2	n=3	n=4
D	-0.0258***	-0.0229**	-0.0377***	-0.0325***	-0.0125**
	(-3.32)	(-2.67)	(-4.22)	(-3.86)	(-2.44)
Control Variables					
Size	-0.0117***	PE	0.0034	Lev	-0.0013
	(-6.61)		(1.48)		(-0.54)
BM	-0.0682***	ROA	-0.0290***	Inv	-0.0144***
	(-20.62)		(-11.28)		(-4.84)
Fixed Effects			✓		
No. Obs			335337		
AR2			0.167		
Panel B: Four-Quarter Operating Income Growth Rates					
	Distance from Shock Origin (in # of Connections)				
	Origin	n=1	n=2	n=3	n=4
D	-0.0543***	-0.0475**	-0.0598***	-0.0543***	-0.0219**
	(-3.46)	(-2.89)	(-3.65)	(-3.18)	(-2.37)
Control Variables			✓		
Fixed Effects			✓		
No. Obs			254322		
AR2			0.106		
Panel C: Four-Quarter Change in Gross Margin					
	Distance from Shock Origin (in # of Connections)				
	Origin	n=1	n=2	n=3	n=4
D	-0.0192*	-0.0154*	-0.0207**	-0.0261**	-0.0108*
	(-2.14)	(-2.05)	(-2.75)	(-2.94)	(-2.12)
Control Variables			✓		
Fixed Effects			✓		
No. Obs			280617		
AR2			0.073		

Table 1.5: Spillover of Idiosyncratic Shocks: Results Scaled by Shocks' Original Impact

This table reports the coefficient estimates of b_n, $n = 0, \ldots, 4$ from Regression (1.2) of the text. b_n measure the incremental impact of the spillover on subsequent connections $n = 1, \ldots, 4$ in units of the percentage impact on the origin firm. The dependent variable is growth rates in revenue. All control variables are defined in Appendix A.1.1. All standard errors are clustered at the firm level. All regressions include industry×year, fiscal quarter, and state/country fixed effects, and are in quarterly frequency from 1994 to 2015.

Four-Quarter Revenue Growth Rates					
	Distance from Shock Origin (in # of Connections)				
	n=1	n=2	n=3	n=4	n=5
D	0.8271***	1.0313***	0.8762***	0.4036**	0.1209
	(3.16)	(3.98)	(3.55)	(2.62)	(1.54)
Control Variables					
Size	-0.0143***	PE	0.0047	Lev	-0.0037
	(-5.61)		(1.32)		(-0.32)
BM	-0.0667***	ROA	-0.0288***	Inv	-0.0129***
	(-18.41)		(-10.93)		(-4.50)
Fixed Effects			✓		
No. Obs			335337		
AR2			0.134		

Table 1.6: Market Power and Shock Spillover, Firm-Level Evidence

This table reports the coefficient estimates of c_n and b_n, $n = 0, \ldots, 4$ from Regression (1.3) of the text. c_n is the first row of each panel and measures the average spillover impact given average values of market power. b_n is the second row and measures the incremental effect of one-standard-deviation change in market power (or market power ratio) on revenue growth rate differences between firms with distance-n shocks and firms without shocks. D^n is a dummy variable that equals to 1 if one of firm i's suppliers from a distance of n connections experiences an idiosyncratic shock captured by the disclosure data. MP is the firm's own market power and MPR is the ratio of the average market power of the firms' suppliers to that of the firm, defined in Section 1.3.3 of the text. Both MP and MPR are standardized to mean of zero and standard deviation of one. All control variables are defined in Appendix A.1.1. All standard errors are clustered at the firm level. All regressions include industry×year, fiscal quarter, and state/country fixed effects, and are in quarterly frequency from 1994 to 2015.

Panel A: Firms' Own Market Power				
	Distance from Shock Origin			
	n=1	n=2	n=3	n=4
D	-0.022**	-0.031**	-0.027**	-0.011*
	(-2.58)	(-2.7)	(-2.94)	(-1.92)
D×MP	0.026***	0.031***	0.024***	0.017***
	(4.73)	(4.86)	(3.45)	(3.61)
Firm Controls	✓	No. Obs	335337	
Fixed Effects	✓	AR2	0.178	

Panel B: Ratio of Own Power to Average Supplier Market Power				
	Distance from Shock Origin			
	n=1	n=2	n=3	n=4
D	-0.020**	-0.028**	-0.025**	-0.009*
	(-2.47)	(-2.81)	(-2.76)	(-2.21)
D×MPR	-0.024***	-0.037***	-0.030**	-0.012***
	(-3.32)	(-3.78)	(-2.94)	(-3.20)
Firm Controls	✓	No. Obs	335337	
Fixed Effects	✓	AR2	0.175	

Table 1.7: Market Power and Shock Spillover, Network-Level Evidence

This table reports the mean, median, and 75th- and 25th-percentile values of MP and MPR measures, defined in Section 1.3.3 of the text, at distances of $n = 0, \ldots, 4$ from the origin of the idiosyncratic shock captured by the disclosure data. The computation uses the lagged value of the $Size$ variable, which is the market capitalization of firms defined in Appendix A.1.1. The power measures are computed as a ratio of firm sizes to total industry sizes at the 4-digit SIC level. All measures are computed in quarterly frequency from 1994 to 2015.

MP	Distance from Shock Origin (in # of Connections)					
	Overall	n=0 (Origin)	n=1	n=2	n=3	n=4
Mean	0.3476	0.4238	0.4868	0.3867	0.2514	0.2292
Median	0.0866	0.0970	0.1109	0.0912	0.0772	0.0752
75th Percentile	0.8834	0.9278	0.9374	0.8723	0.8655	0.8528
25th Percentile	0.0249	0.0460	0.0455	0.0246	0.0105	0.0136
No. Obs	335337	8160	40469	36477	45491	44290

Table 1.8: Input Substitutability, Inventories, and Shock Spillover

This table reports the coefficient estimates of c_n and b_n, $n = 0, \ldots, 4$ from Regressions (1.5a) and (1.5b) of the text. c_n is the first row of each panel and measures the average spillover impact given average values of market power. b_n is the second row and measures the incremental effect of one-standard-deviation change in inventory (or supplier substitutability) on revenue growth rate differences between firms with distance-n shocks and firms without shocks. D^n is a dummy variable that equals to 1 if one of firm i's suppliers from a distance of n connections experiences an idiosyncratic shock captured by the disclosure data. $INVR$ is the inventory-to-total-assets level and $\overline{\gamma}$ is average supplier share, defined in Section 1.3.3 of the text. Both $INVR$ and $\overline{\gamma}$ are standardized to mean of zero and standard deviation of one. The supplier substitutability measure uses a reduced sample where the specific values are available for each link. All control variables are defined in Appendix A.1.1. All standard errors are clustered at the firm level. All regressions include industry×year, fiscal quarter, and state/country fixed effects, and are in quarterly frequency from 1994 to 2015.

Panel A: Inventory Levels				
	Distance from Shock Origin (in # of Connections)			
	n=1	n=2	n=3	n=4
D	-0.023**	-0.039***	-0.033***	-0.011**
	(-2.82)	(-3.17)	(-3.24)	(-2.69)
D×INVR	0.010***	0.011***	0.013***	0.009***
	(4.36)	(4.52)	(4.39)	(4.45)
Firm Controls	✓	No. Obs	335337	
Fixed Effects	✓	AR2	0.171	

Panel B: Supplier Substitutability (Sample with $\gamma < 1$ only)				
	Distance from Shock Origin			
	n=1	n=2	n=3	n=4
D	-0.020**	-0.023**	-0.023**	-0.009*
	(-2.51)	(-2.74)	(-2.59)	(-1.96)
D×$\overline{\gamma}$	0.035***	0.052***	0.036***	0.031***
	(4.80)	(4.84)	(3.31)	(3.11)
Firm Controls	✓	No. Obs	249926	
Fixed Effects	✓	AR2	0.189	

Table 1.9: Robustness: Ensuring Shocks Have Only Firm-Specific Effects

This table reports the coefficient estimates of b_n, $n = 0, \ldots, 4$ from Regression (1.1) of the text, on 11 reduced samples where I either remove one shock category at a time (Panel A), or use one shock category at a time (Panel B). The shock categories are classified by the LDA algorithm and defined in Section 1.2.1 of the text. The "Fire Only" category is a subset of shocks from the "Manmade" category that pertains to localized fires only, and is constructed according to Section 1.4.1 of the text. b_n measures the average difference between firms hit with a shock spilled over from a distance of n connections, and firms never hit with any shocks. D^n is a dummy variable that equals to 1 if one of firm i's suppliers from a distance of n connections experiences an idiosyncratic shock captured by the disclosure data. All control variables are defined in Appendix A.1.1. All standard errors are clustered at the firm level. All regressions include industry×year, fiscal quarter, and state/country fixed effects, and are in quarterly frequency from 1994 to 2015.

Panel A: Remove Individual Shock Categories

	Category Removed					
	(0)	(1)	(2)	(3)	(4)	(5)
	None	Disaster	Manmade	Breakdown	IT	Adjustment
Origin Firms	-0.0258***	-0.0252***	-0.0231***	-0.0225***	-0.0286***	-0.0292***
	(-3.32)	(-3.18)	(-3.74)	(-3.22)	(-3.94)	(-3.47)
Distance 1 Firms	-0.0229**	-0.0241**	-0.0210**	-0.0204**	-0.0251**	-0.0224**
	(-2.67)	(-2.96)	(-2.60)	(-2.62)	(-3.03)	(-2.71)
Firm Controls	✓	✓	✓	✓	✓	✓
Fixed Effects	✓	✓	✓	✓	✓	✓
No. Obs	335337	335337	335337	335337	335337	335337
AR2	0.167	0.153	0.162	0.155	0.166	0.159

Panel B: Use Individual Shock Categories and Fire Only

	Category Used					
	(1)	(2)	(3)	(4)	(5)	(6)
	Fire Only	Disaster	Manmade	Breakdown	IT	Adjustment
Origin Firms	-0.0174**	-0.0247***	-0.0275**	-0.0288***	-0.0191*	-0.0199**
	(-2.87)	(-3.66)	(-2.73)	(-3.96)	(-2.03)	(-2.84)
Firm Controls	✓	✓	✓	✓	✓	✓
Fixed Effects	✓	✓	✓	✓	✓	✓
No. Obs	335337	335337	335337	335337	335337	335337
AR2	0.109	0.134	0.138	0.145	0.120	0.117

Table 1.10: Robustness: Prior Growth Trends

This table reports the coefficient estimates of b_n, $n = 0, \ldots, 4$ from Regression (1.1) of the text. b_n measures the average difference between firms hit with a shock spilled over from a distance of n connections, and firms never hit with any shocks. The dependent variables are lagged growth rates in revenue from the previous 1, 2, 4, and 8 quarters prior to the shocks' quarter. D^n is a dummy variable that equals to 1 if one of firm i's suppliers from a distance of n connections experiences an idiosyncratic shock captured by the disclosure data. All control variables are defined in Appendix A.1.1. All standard errors are clustered at the firm level. All regressions include industry×year, fiscal quarter, and state/country fixed effects, and are in quarterly frequency from 1994 to 2015.

Prior revenue growth trends					
	Distance from Shock Origin (in # of Connections)				
	Origin	n=1	n=2	n=3	n=4
t-1→t	0.0012	-0.0004	-0.0004	-0.0013	0.0003
	(0.72)	(-0.54)	(-0.83)	(-0.61)	(0.49)
t-2→t	-0.0030	-0.0033	0.0009	-0.0016	-0.0019
	(-1.25)	(-0.89)	(1.43)	(-1.31)	(-0.62)
t-4→t	-0.0036*	0.0076	0.0039	0.0008	-0.0026
	(-1.67)	(1.53)	(1.22)	(0.69)	(-1.08)
t-8→t	0.0106	-0.0056	0.0097	0.0103	-0.0034
	(0.75)	(-0.41)	(1.19)	(0.58)	(-0.87)
Firm Controls	✓	✓	✓	✓	✓
Fixed Effects	✓	✓	✓	✓	✓

Table 1.11: Robustness: Falsification Test with Random Shocks

This table reports the coefficient estimates of b_n, $n = 0, \ldots, 4$ from Regression (1.1) of the text. b_n measure the incremental impact of the spillover on subsequent connections $n = 1, \ldots, 4$ in units of the percentage impact on the origin firm. The dependent variable is growth rates in revenue. The first row of the independent variables is the real shocks: D^n is a dummy variable that equals to 1 if one of firm i's suppliers from a distance of n connections experiences an idiosyncratic shock captured by the actual disclosure data. $FAKED$ is the falsified shocks: they are shocks randomly given to other firms at the time of the real shocks, constructed according to Section 1.4.2 of the text. All control variables are defined in Appendix A.1.1. All standard errors are clustered at the firm level. All regressions include industry×year, fiscal quarter, and state/country fixed effects, and are in quarterly frequency from 1994 to 2015.

Four-Quarter Revenue Growth Rates					
	Distance from Shock Origin (in # of Connections)				
	Origin	n=1	n=2	n=3	n=4
Real Shocks	-0.0258***	-0.0229**	-0.0377***	-0.0325***	-0.0125**
	(-3.32)	(-2.67)	(-4.22)	(-3.86)	(-2.44)
Fake Shocks	0.0057	0.0102	0.0024	-0.0058	0.0035
	(1.02)	(0.79)	(1.23)	(-0.55)	(0.64)
Firm Controls	✓	✓	✓	✓	✓
Fixed Effects	✓	✓	✓	✓	✓

Table 1.12: Robustness: Strategic Disclosures

This table reports the coefficient estimates of b_n, $n = 0, \ldots, 4$ from Regression (1.1) of the text, on a series of subsamples. b_n measures the average difference between firms hit with a shock spilled over from a distance of n connections, and firms never hit with any shocks. Subsamples A and B includes all observations prior to, and after, August 23, 2004, respectively, corresponding to the enforcement date of Provision 209 of the Sarbanes-Oxley Act. Subsample C consists of internal shocks disclosed by the firm. See Section 1.4.1 of the text for details. D^n is a dummy variable that equals to 1 if one of firm i's suppliers from a distance of n connections experiences an idiosyncratic shock captured by the disclosure data. All control variables are defined in Appendix A.1.1. All standard errors are clustered at the firm level. All regressions include industry×year, fiscal quarter, and state/country fixed effects, and are in quarterly frequency from 1994 to 2015.

Coefficient for shock dummy using different subsamples					
	Distance from Shock Origin (in # of Connections)				
	Origin	n=1	n=2	n=3	n=4
Full Sample	-0.0258***	-0.0229**	-0.0377***	-0.0325***	-0.0125**
	(-3.32)	(-2.67)	(-4.22)	(-3.86)	(-2.44)
Subsamples					
A: Pre-SOX Sample	-0.0206***	-0.0249**	-0.0350***	-0.0338***	-0.0119**
	(-3.28)	(-2.61)	(-4.17)	(-3.54)	(-2.46)
B: Post-SOX Sample	-0.0272***	-0.0215**	-0.0387***	-0.0324***	-0.0130**
	(-3.36)	(-2.77)	(-4.10)	(-3.85)	(-2.79)
C: Excluded External Disclosures	-0.0279***	-0.0268**	-0.0372***	-0.0339***	-0.0117**
	(-3.44)	(-2.95)	(-4.07)	(-3.91)	(-2.36)
Firm Controls	✓	✓	✓	✓	✓
Fixed Effects	✓	✓	✓	✓	✓

Table 1.13: Corporate Response to Shock Spillovers: Capital Buildup

This table reports the coefficient estimates of b_n, $n = 0, 1, > 1$, from Regression (1.6) of the text. b_n measures the average difference in corporate policies between firms hit with a shock spilled over from a distance of n connections, and firms never hit with any shocks. The dependent variables are changes in cash (CHEQ), inventory (INVTQ), capital expenditures (CAPXQ), and R&D expenditures (XRDQ), scaled by lagged total assets (ATQ), from the previous quarter to the quarter of shocks, and from the shock quarters to the subsequent 1, 4, and 8 quarters. D^n is a dummy variable that equals to 1 if one of firm i's suppliers from a distance of n connections experiences an idiosyncratic shock captured by the disclosure data. All dependent variables are standardized to mean of zero and standard deviation of one. All control variables are defined in Appendix A.1.1. All standard errors are clustered at the firm level. All regressions include industry×year, fiscal quarter, and state/country fixed effects, and are in quarterly frequency from 1994 to 2015.

Change in Capital Expenditures												
	Working Capital						Other Capital Expenditures					
	Inventory			Cash			CAPEX			R&D		
	Origin	n=1	n¿1	Origin	n=1	n¿1	Origin	n=1	n¿1	Origin	n=1	n¿1
t-1,t	-0.001	-0.003	-0.005	-0.001	0.001	0.000	-0.005	0.005	-0.002	0.016	-0.001	0.003
	(-0.40)	(-1.09)	(-0.82)	(-1.01)	(1.24)	(0.25)	(-1.16)	(0.84)	(-0.41)	(-0.70)	(-0.56)	(1.02)
t,t+1	-0.083***	-0.072***	-0.067***	-0.020***	-0.022***	-0.020*	0.001	0.002	0.010	-0.003	0.000	-0.004
	(-4.67)	(-5.31)	(-5.58)	(-3.87)	(-3.31)	(-2.04)	(0.84)	(1.53)	(1.12)	(-0.83)	(0.48)	(-0.63)
t,t+4	0.095***	0.090***	0.003	-0.023*	-0.016	0.002	0.030**	0.026**	-0.002	0.006	0.009	-0.001
	(4.85)	(5.79)	(1.12)	(-2.20)	(-1.00)	(1.13)	(2.35)	(2.90)	(-0.32)	(1.12)	(1.04)	(-0.31)
t,t+8	0.148**	0.100**	0.003	0.047**	0.063**	0.001	0.045**	0.058**	0.015*	0.042	0.034*	0.003
	(3.03)	(2.96)	(0.27)	(2.56)	(2.78)	(0.84)	(2.59)	(2.66)	(2.07)	(1.48)	(2.19)	(1.10)

Table 1.14: Corporate Response to Shock Spillovers: Financing Policies

This table reports the coefficient estimates of b_n, $n = 0, 1, > 1$, from Regression (1.6) of the text. b_n measures the average difference in corporate policies between firms hit with a shock spilled over from a distance of n connections, and firms never hit with any shocks. The dependent variables are changes in leverage (DLTT-DLC divided by market capitalization), net debt issuance (DLTISQ-DLTRQ), net equity issuance (SSTKQ-PRSTKCQ), account payables (APQ), and retained earnings (REQ), scaled by lagged total assets (ATQ), from the previous quarter to the quarter of shocks, and from the shock quarters to the subsequent 1, 4, and 8 quarters. D^n is a dummy variable that equals to 1 if one of firm i's suppliers from a distance of n connections experiences an idiosyncratic shock captured by the disclosure data. All dependent variables are standardized to mean of zero and standard deviation of one. All control variables are defined in Appendix A.1.1. All standard errors are clustered at the firm level. All regressions include industry×year, fiscal quarter, and state/country fixed effects, and are in quarterly frequency from 1994 to 2015.

Change in Capital Structure									
	Leverage			Financing 1					
	Market Leverage			Net Debt Issue			Net Equity Issue		
	Origin	n=1	n¿1	Origin	n=1	n¿1	Origin	n=1	n¿1
t-1,t	0.007	-0.002	0.006	0.012	-0.002	0.003	-0.001	-0.005	0.011
	(0.85)	(-0.14)	(0.97)	(1.54)	(-0.31)	(0.56)	(-0.15)	(-0.98)	(0.40)
t,t+1	0.015	-0.023	0.000	-0.009	0.004	-0.004	-0.004	0.007	0.005
	(1.34)	(-1.48)	(0.62)	(-0.38)	(0.57)	(-1.29)	(-0.81)	(1.14)	(0.61)
t,t+4	0.086***	0.079*	0.008	0.028*	0.035**	-0.009	-0.002	-0.006	-0.001
	(3.18)	(1.94)	(0.60)	(2.22)	(2.98)	(-0.10)	(-0.29)	(-1.32)	(-0.30)
t,t+8	0.085**	0.094***	0.022**	0.047***	0.043***	0.014*	0.011	-0.003	-0.002
	(3.07)	(3.36)	(2.42)	(3.59)	(3.87)	(2.13)	(1.05)	(-1.41)	(-0.77)

	Financing 2					
	Retained Earnings			Trade Credits (Payables)		
	Origin	n=1	n¿1	Origin	n=1	n¿1
t-1,t	-0.002	-0.011	-0.015	0.001	-0.006	-0.005
	(-0.62)	(-1.25)	(-1.09)	(0.21)	(-0.83)	(-0.80)
t,t+1	-0.040*	-0.066*	-0.010	0.000	-0.007	-0.002
	(-2.21)	(-1.88)	(-1.20)	(0.13)	(-0.74)	(-0.47)
t,t+4	-0.025**	-0.017	-0.024*	-0.008	-0.020*	0.002
	(-2.48)	(-1.49)	(-2.22)	(-1.39)	(-1.87)	(0.94)
t,t+8	0.007	-0.013	-0.006	0.003	-0.012	-0.004
	(1.41)	(-1.50)	(-0.49)	(0.14)	(-1.55)	(-1.02)

CHAPTER 2 : Deciphering Fedspeak: The Information Content of FOMC Meetings

(with Narasimhan Jegadeesh)

2.1. Introduction

Monetary policies implemented by the Federal Reserve have profound effects on the global economy. Numerous papers in the economics and finance literature examine the determinants and effects of such policies using *quantitative* "Fed proxies" such as the federal funds target rate or the reserve requirement. In addition to these hard data, the Fed routinely releases large amounts of *qualitative* information, such as meeting minutes, transcripts, and speeches, in an effort to foster effective communication with the public and achieve greater operational transparency. While a voluminous literature examines market reactions to quantitative information such as rate changes, very few papers explore the informativeness of these "soft" data conveyed in the language of Fed communications. Do they have incremental information value? How does the market react to these data? Can they be used as alternative predictors of economic and policy outcomes?

Our paper fills the void by presenting an innovative, topic-based approach to determine the informativeness of FOMC meeting minutes, which are detailed summaries of everything discussed during the preceding meeting. Because such discussions encompass a wide range of topics, the proportions of which vary widely from meeting to meeting, we use an automated algorithm based on Bayesian learning to objectively and robustly classify each individual paragraph in the minutes into four distinct economic themes that intuitively correspond to specific Fed mandates and tasks: *Growth*, *Inflation*, *Financial markets*, and *Policy*. We then simultaneously extract contents–the tone and uncertainty level–from the texts of each minutes, and by topic. Compared to a manual approach such as Romer and Romer (1989), our objective approach minimizes any potential researcher-induced bias, thereby allowing us to accurately gauge the specific context of each discussion and, for each meeting minutes, obtain a granular measure of topic mix that human readers cannot accurately identify.

To further remove any subjectivity, we assess the informativeness of each topic based on financial market's reactions to the release of the minutes.

We find several new results with our approach. First, we demonstrate that the texts of FOMC minutes contain incremental information not incorporated in either rate announcements or the more timely meeting statements, despite the fact that the minutes are released several weeks after the meetings. Lucca and Moench (2015) find strong evidence that policy announcements on the day of the FOMC meeting is associated with significantly higher stock market volatilities both on and prior to the meeting days. We show that, several weeks after the meeting dates, the release of the minutes is also correlated with a similar degree of volatility spike in both equity and debt markets.

Our next set of tests examine the granular source of this additional informativeness from individual topics. We first demonstrate that, when treated as a single unit, each document as a whole does not yield informative results: neither whole-document tone nor uncertainty is significantly related to market reaction. However, the market do find the discussion on individual topics informative, and assign different informational value to different topics. The market finds traditional "dual mandate" themes, such as *Inflation*, most informative. Interestingly, the market also reacts strongly to the content of the relatively new topic of *Financial markets*, reflecting the Fed's increasingly important role of maintaining systemic stability, particularly during and after the recent financial crisis. Furthermore, we find that the *Policy* topic is not only deemed informative by the market, but its discussion is also orthogonal to existing economic conditions, indicating that the FOMC members do not necessarily follow fixed guidelines such as the Taylor rule when setting the monetary policy. Our topic-content Scores also hold significant predictive power for real economic activities, which we explore in a related research.

The results above suggest that the Fed possesses superior information than other market participants. Our next tests examine whether such superior information is transmitted to the market through "soft channels" conveyed by language of the minutes. We show that the price jumps at the release of the minutes do not revert, and market volatility is greatly reduced after the release of the minutes. This is consistent with information transmission into the market at the time of the minutes' release.

Our paper contributes to the literature on three fronts. First, our paper is the first in finance to use a topic-based textual analysis approach on the FOMC minutes, and our approach provides collection of intuitive indicators on multiple facets of the economy and monetary policy, which are also orthogonal from existing economic variables. Alternative text-based economic indicators also exist, such as Baker et al. (2015), which is based on counting the frequency of uncertainty-related words in news reports. By contrast, our policy indicators are derived directly from the language of policy makers themselves. Unlike news reports, each FOMC minutes is likely to be painstakingly scrutinized by the market, and the usage of every word from the minutes thus matters. This is evidenced by the significant market reaction to our measures. As such, our economic and policy indicators are likely to contain more policy-relevant information and less noise.

Second, our paper furthers the burgeoning literature of financial textual analysis by being the first to employ a paragraph-level information retrieval system that moves beyond the traditional word-based approach employed in current literature such as Tetlock (2007), Hanley and Hoberg (2010), Loughran and McDonald (2011), and Jegadeesh and Wu (2013). This paper is the first in finance to employ on FOMC minutes the Latent Dirichlet Allocation (LDA) model of automated topic retrieval, which has been successfully employed to characterize topics of a wide variety of document sources, from journal articles in Nature to patient-discharge reports.[1] Compared to word-based alternative approaches such as Singular Value Decomposition used by Boukus and Rosenberg (2006), Bayesian methods that explicitly account for the distribution of both topics and words such as the LDA are ideally suited to our collection of FOMC minutes for the following reasons: first, the topic mix and content of FOMC minutes are sufficiently varied, which leads to both robust and intuitively

[1] For a list of LDA applications and an evaluation of their effectiveness, see Blei et al. (2003).

appealing classification results that are on par with or exceeds manual classification by researchers.[2] Second, compared to manual approaches, our approach is entirely objective, relying only on the structure of the provided texts, and does not require subjective input from researchers. Third, many paragraphs in the FOMC minutes exhibit several topics without a dominant topic. In this case researchers would have difficulty manually identifying the proportion of each topic, while our algorithm outputs the proportion directly, enabling us to compute a unified topic-content score for each minutes.

Furthermore, we provide a model-free alternative of time-varying monetary policy. Structural models such as Ang et al. (2002), Campbell et al. (2015) and Sims and Zha (2006) usually posit the existence of latent policy "regimes" beyond the observable data, and estimate such regimes in a structural VAR setting. However, the specific mechanism from which policies are generated depends on the underlying model supplied by the researcher, which can be subjective. By contrast, our approach directly outputs the economic and policy contents from the texts of FOMC minutes. Our Policy Score series can be interacted with any identifiable economic variables, thereby explicitly generating "latent" states such as policy *tone*, *aggressiveness*, or *uncertainty*, etc. Therefore, our text-based measure nicely complements the interest-rate-based structural models by providing additional rich data moments.

The rest of the paper is organized as follows. Section 2.2 describes our sample and data sources. Section 2.3 introduces our automated, topic-based content analysis methodology. Section 2.4 reports the results of our empirical tests and explores the sources of predictive power of our measures. Section 2.6 concludes.

[2] We manually select 50 paragraphs and employ 10 research assistants to classify them manually into our topic collection and to identify the topic mixture. On average the algorithm agrees with human researchers in 46 out of 50 cases. See Section 2.3.2 for details.

2.2. Data

2.2.1. Introduction of FOMC Meetings and Minutes

This subsection provides a brief overview of the logistic details of FOMC meetings and the release of the meeting minutes. From the early 1980s, the FOMC holds eight regularly scheduled meetings per year, during which members discuss the economic outlook and formulate monetary policy. Any policy change decided at the meeting is immediately implemented through open market operations. Prior to 1994, no public announcement about policy was made and the market inferred any policy change through the size and direction of the open market operations on the next day. Starting from January 1994, specific policy changes were made public in a short *meeting statement* released immediately after the meeting.

Moreover, during each meeting, detailed records of the discussions are kept, then summarized in the form of *meeting minutes*, which are released to the public after a delay.[3] The minutes contain no new information received between the meeting date and the release date, and instead serve as an overview of the members' internal discussions on their economic outlook, as well as a nuanced explanation of the rationale for any policy change.

The meeting minutes follow a highly structured writing style. They are routinely consisted of four major sections. The first section outlines the administrative detail of the meeting and reviews previous open market operations. The second section provides the staff's review and outlook of the economic and financial situation, prepared in advance of the meeting. The next two sections provide the bulk of the economic content: the third section details the FOMC members' discussion of the current economic and financial situation, as well as their own economic outlook and projections. The last section is mostly related to policy and discusses the rationale for current policy and outlook for future policies. We remove the first section prior to processing the documents since it is unlikely to contain any economically

[3]The delay ranges between three and eight weeks. The Fed implemented a series of accelerated release schedule during the 1990s and 2000s, which shortened the lag from eight (before 2004) from three weeks (after 2004). From 1997 onward, the minutes are released at 2:00pm Eastern Standard Time.

meaningful content.

2.2.2. The FOMC Minutes Sample

We download all FOMC meeting minutes between the February 1991 and June 2015 meetings from FOMC's web site. Some minutes in earlier periods are only available in scanned PDF format, and we and obtain all textual data from these PDF documents using a text extraction engine.[4] We also record the date of the meeting, and the date and earliest time of the release of each minutes by examining the timestamp of the released file. Our sample consists of 196 meeting minutes (thereafter referred to as Minutes).

For each Minutes, we develop a textual parsing algorithm to simultaneously achieve the following: 1) remove the introductory section of the Minutes that lists participant names and administrative matters, and remove the section on specific open market operations (e.g. amount of securities purchased); 2) break the document into individual paragraphs; 3) record the specific section where each paragraph is located (e.g. Staff Economic Discussion or Members' Discussion), and, 5) obtain paragraph length in the number of words. This procedure produces 28,676 unique sentences and 5,644 paragraphs. The average sentence length is 29 words.

2.2.3. Market Reaction Data

In many of our tests, we use high-frequency trading data from both equity and bond markets in order to measure market reactions to the contents of the minutes as broadly as possible. For the equity market, our main instrument is the tick-by-tick trading data from the SPDR exchange-traded fund by State Street to proxy for the overall level of stock market response. The SPDR, launched in 1993, follows the S&P 500 index with negligible tracking error. Trading volume has increased dramatically since 2000, making SPDR one of the most liquid stocks. Since volume prior to 2000 is low, we restrict our sample period from 2000 to 2015. As an additional robustness check, we also use proprietary data on the S&P E-

[4]Minutes downloaded in PDF at http://www.federalreserve.gov/monetarypolicy/fomccalendars.htm

Mini futures contracts from the Chicago Mercantile Exchange (CME), which offers similar liquidity levels post-2000. Our results are similar using both instruments.

For the bond markets, we use high-frequency electronic trading data for the Eurodollar futures contracts obtained from the CME. To construct the trading history, we use the "front month" contract, which is the one with expiration date closest to the trading date. Electronic trading was sporadic prior to 2003, and as a result, for Eurodollar futures, we can only construct a reliable trading history for a shorter sample period from 2003 to 2014.

Next, we construct our event window around the time when the meeting minutes are released. We then calculate return volatility during the event window. After 1997, the official release time for the meeting minutes is 2:00pm Eastern Standard Time. However, it is possible that some minutes are released early or late. As such, for each release day, we use an automated algorithm to simultaneously search the FOMC's official web site, Bloomberg, Dow Jones Newswires, and Thomson Reuters, and comparing the time on the FOMC timestamp with that of the first news story of the same day on the Minutes' release. We record the release time as the the earliest time that the minutes (or news about the minutes) are reported among the these sources. The actual release time ranges between 1:59pm and 2:06pm. Therefore, we construct our event window as the 15-minute window between 2:00pm and 2:15pm each day. Our result is robust to alternative event window specifications. The results are also similar using windows ranging from 20 minutes to two hours.[5]

We then calculate event-window return and, following convention, raw return volatility for the equity market is computed as the squared event-window return and that for the Eurodollar market is computed as the absolute value of yield changes. Specifically, for each minutes t in our sample,

$$R_t^{SPY} = \frac{P_{t,2:15pm}^{SPY} - P_{t,2:00pm}^{SPY}}{P_{t,2:00pm}^{SPY}} \tag{2.1a}$$

[5] We have used windows starting as early as 1:50pm to as late as 2:05pm. We also used window lengths from 20 minutes to 2 hours, in 10-minute increments. The results are similar thoughout most window lengths.

$$R_t^{ED} = \frac{Y_{t,2:15pm}^{ED} - Y_{t,2:00pm}^{ED}}{Y_{t,2:00pm}^{ED}} \tag{2.1b}$$

$$V_t^{SPY} = \left(R_t^{SPY}\right)^2 = \left(\frac{P_{t,2:15pm}^{SPY} - P_{t,2:00pm}^{SPY}}{P_{t,2:00pm}^{SPY}}\right)^2 \tag{2.1c}$$

$$V_t^{ED} = \left|R_t^{ED}\right| = \left|\frac{Y_{t,2:15pm}^{ED} - Y_{t,2:00pm}^{ED}}{Y_{t,2:00pm}^{ED}}\right| \tag{2.1d}$$

Because we use a very short, 15-minute window in constructing the market volatility measure, confounding effects from other macroeconomic variables are negligible, as the minutes are released predominantly on Wednesdays and (before 2004) Thursdays, and no other significant economic indicators are released on these afternoons.[6] To further ensure that any volatility change during our short event window is solely a contemporaneous response to the minutes' release, rather than a delayed response to other macroeconomic events, we separate the event window volatility into an expected and unexpected part. Specifically, we compute the unexpected volatility on the release day as the difference between the raw volatility and the average event window volatility, computed per Equations (2.1c) and (2.1d), in the past k trading days:

$$UV_{t,k} = V_t - \sum_{j=1}^{k} \frac{V_{t-j}}{k} \tag{2.2}$$

In general we set k between 5 and 30 trading days. Most results in our Tables are reported using $k = 20$ days. The results are little changed when k is set to other lengths. We therefore omit the k-subscript and instead use the notation in UV_t subsequent discussions.

2.3. Methodology

Because each Minutes is a summary of everything that is discussed during the preceding meeting, it is a mixture of a wide range of topics. This is demonstrated by several excerpts from the minutes that we present in Appendix A.2: while one paragraph discusses the

[6] See http://www.bloomberg.com/markets/economic-calendar for a schedule of important economic news. Usually no other significant news are scheduled to release on Wednesdays. On Thursdays most other indicators are released on the morning prior to market open.

latest developments on inflation, another paragraph might provide outlook on financial markets. Another paragraph might discuss both. Several complications arise from these multi-faceted texts: First, which discussions are informative and which are not? Second, many words have different connotations under different contexts. For example, *increase* is considered a positive word in the economic growth context, but has negative connotations in the inflation context. How do we separate these contexts? Third, the proportion and content of discussions on each topic are likely to vary from meeting to meeting. How should one accurately measure these proportions?

These are our motivation for using a topic-based approach that isolates the content of each topic prior to content extraction. This approach allows us to address the above concerns simultaneously by 1) on the paragraph level, accurately gauging the context of each paragraph, and 2) on the document level, obtaining a granular measure of time-varying content proportions that human readers cannot accurately identify. Overall, our approach adds another dimension that enhances traditional content analysis. This section describes our methodology to separate the FOMC minutes into individual topics and extract the content from each topic.

2.3.1. The Latent Dirichlet Allocation (LDA) Algorithm

We first classify each Minutes into distinct topics with the Latent Dirichlet Allocation (LDA) algorithm first developed by Blei et al. (2003), which belongs to a broader class of probabilistic topic models that use hierarchical Bayesian analysis to uncover the underlying semantic structure of textual documents. The common intuition behind such topic models can be summarized by two statistical distributions, which constitute the latent data generating process: The base unit of our analysis is a paragraph. Each paragraph is sufficiently summarized as a distribution over a collection of topics, each of which is, in turn, a distribution over the collection of English words used in the sample texts. For example, a paragraph that discusses inflation should be represented by a distribution that places a high weight on a topic that places high weights on words such as *prices, CPI, inflation*, etc.

By contrast, a topic that places high weights on *foreign trade* and *imports* should receive a low weight in this paragraph distribution.

However, the two distributions are unobservable from the point of the researcher. The advantage of probabilistic topic models is that, using Bayesian techniques, such models efficiently infer the hidden distributional properties from the observable data (i.e. the collection of documents). LDA represents one particular parameterization of the model: We assume that these two latent distributions belong to the Dirichlet family. Then, armed with this functional form and the observed words in each paragraph, we compute the posterior (i.e. empirical) paragraph and topic distributions using the standard Bayes Theorem. These empirical distributions are the main outputs of the model. The only inputs in LDA are the document texts and the number of topics. As such, compared to a manual classification approach, researcher-induced subjectivity and bias are minimized.

We illustrate our approach with a simple example. Suppose that the full set of revelant FOMC vocabulary consists of only $V = 4$ words (ignore common words such as *we*, *the*, etc): {*employment*, *layoff*, *imports*, *trade*}. We are given $D = 3$ paragraphs:

1. *Employment situation is good and layoff has declined.*

2. *Imports have increased and the outlook for trade is good.*

3. *Imports look good, and employment situation is also good.*

A human reader would intuitively recognize that the first paragraph is about employment and the second is about foreign trade. The third paragraph is a mixture of both. Suppose we fit the LDA model with $N = 2$ topics. If the model performs satisfactorily, then first, the posterior topic distributions should clearly and intuitively identify the topics and thus be something similar to:

- $\hat{\beta}_1 \equiv \{\hat{P}_{topic1}(employment), \hat{P}_{topic1}(layoff), \hat{P}_{topic1}(imports), \hat{P}_{topic1}(trade)\}$

  $= \{0.55, 0.43, 0.01, 0.01\}$

- $\hat{\beta}_2 \equiv \{\hat{P}_{topic2}(employment), \hat{P}_{topic2}(layoff), \hat{P}_{topic2}(imports), \hat{P}_{topic2}(trade)\}$

  $= \{0.01, 0.01, 0.60, 0.48\}$

Next, the posterior topic mixture in each paragraph should correspond to the human reader's intuition:

- $\hat{\theta}_1 \equiv \{\hat{P}_{paragraph1}(Topic1), \hat{P}_{paragraph1}(Topic2)\} = \{0.99, 0.01\}$
- $\hat{\theta}_2 \equiv \{\hat{P}_{paragraph2}(Topic1), \hat{P}_{paragraph2}(Topic2)\} = \{0.01, 0.99\}$
- $\hat{\theta}_3 \equiv \{\hat{P}_{paragraph3}(Topic1), \hat{P}_{paragraph3}(Topic2)\} = \{0.51, 0.49\}$

We proceed with our LDA classification of the FOMC minutes simply by generalizing this example to our sample of $D = 5,644$ unique paragraphs. This set of paragraphs becomes our document collection and our input to the LDA algorithm. Stop words, such as *a, the*, etc., are removed prior to processing. This results in a collection of $V = 61,432$ words.

Next, we hypothesize that there are $N = 8$ unique topics in the document. Our results are robust to alternative specifications from $N = 5$ to $N = 10$.[7] This is the only manual step in the entire process. Here, each of the N topics represents a distribution over the V words in the FOMC vocabulary, and each paragraph is a mixture of the N topics. We assume that the observable data, i.e. words in each document, is generated from a probabilistic data generating process parameterized as follows:

1. Each of paragraph $d = 1, \ldots, D$ contains a mixture of N topics. Let the proportion of topic n in paragraph d be $\theta_{d,n}$ and let the vector $\theta_d = [\theta_{d,1}, \ldots, \theta_{d,N}]'$ represent the true topic mixture of paragraph d. For each d, we assume that this mixture follows an order-N Dirichlet distribution over the N topics, governed by the latent, parameter vector μ of size N:

$$\theta_d \sim Dirichlet_N(\mu)$$

[7] Because each FOMC minutes contains at least four sections, it is likely that $N \geq 5$. When the number of topics increase, some topics become redundant. However, the algorithm results in a similar number of major topics after grouping similar topics as discussed below.

2. Given paragraph d's topic mixture θ_d, let the assignment of each word i in document d into topics be $Z_{d,i}$, where $Z_{d,i} \in \{1, \ldots, N\}$. We assume that this assignment follows the multinomial distribution governed by the document-specific topic vector θ_d described in the previous step:

$$Z_{d,i}|\theta_d \sim Multinomial(\theta_d) \tag{2.3}$$

   Suppose there are I_d unique words in document d. Let the vector Z_d denote the collection of the topic assignment of all words within d, i.e. $Z_d = \{Z_{d,i}\}_{i=1}^{I_d}$

3. The N topic distributions (applied universally to all paragraphs) are in the collection $\beta = \{\beta_1, \ldots, \beta_N\}$. Each topic β_n also follows an order-V Dirichlet distribution over the V words, governed by the latent scalar parameter ϕ:

$$\beta_n \sim Dirichlet_V(\phi) \tag{2.4}$$

4. For each word i in document d, there are V choices to choose from based on our FOMC vocabulary. Conditional on the chosen topic for word i in Step 2 above (i.e. a draw from Distribution (2.3)), and on the structure of the topic distribution from Step 3 (i.e. a draw from Distribution (2.4)), we assume that actual choice of the word, $W_{d,i}$, follows a multinomial distribution governed by the resulting word-topic assignment $\beta_{Z_{d,i}}$:

$$W_{d,i}|\left(\{\beta_n\}_{n=1}^{N}, Z_{d,i}\right) \sim Multinomial(\beta_{Z_{d,i}})$$

   Similarly, let the W_d denote the collection of the vocabulary choice of all words within document d: $W_d = \{W_{d,i}\}_{i=1}^{I_d}$

The above four distributions constitute the latent data generating process that results in our observable document collection $\{W_d\}_{d=1}^{D}$. Recall that they are not directly observable to the researcher. Instead, the only observable data is the occurrence of the actual words i in each document d, i.e. W_d. We can then write the overall data generating process as the joint distribution of latent variables $\{\beta_n\}_{n=1}^N, \{\theta_d\}_{d=1}^D, \{Z_d\}_{d=1}^D$ and the observable variable $\{W_d\}_{d=1}^D$:

$$\begin{aligned}
&P\left(\{\beta_n\}_{n=1}^N, \{\theta_d\}_{d=1}^D, \{Z_d\}_{d=1}^D, \{W_d\}_{d=1}^D\right) \\
&= \prod_{n=1}^N P(\beta_n) \prod_{d=1}^D P(\theta_d) \left[\prod_{i=1}^{I_d} P(Z_{d,i}|\theta_d) P\left(W_{d,i}|\{\beta_n\}_{n=1}^N, Z_{d,i}\right)\right]
\end{aligned}$$

where $P(\cdot)$ are the respective (Dirichlet or multinomial) density functions specified above.

Now that we observe our FOMC document collection $\{W_d\}_{d=1}^D$, we can compute the posterior distribution of the document-topic structure given the observed documents using Bayes' Rule:

$$P\left(\{\beta_n\}_{n=1}^N, \{\theta_d\}_{d=1}^D, \{Z_d\}_{d=1}^D | \{W_d\}_{d=1}^D\right) = \frac{P\left(\{\beta_n\}_{n=1}^N, \{\theta_d\}_{d=1}^D, \{Z_d\}_{d=1}^D, \{W_d\}_{d=1}^D\right)}{P\left(\{W_d\}_{d=1}^D\right)}. \tag{2.5}$$

Similar to other Bayesian inference methods, the numerator in Equation (2.5) and can be easily computed. The denominator is by construction a double integral and therefore cannot be feasibly computed. However, it can be efficiently approximated using a Gibbs sampler. We use a customized Gibbs sampler for fast implementation and defer the technical aspects of the Bayesian inference to the Online Appendix.

2.3.2. Results from the LDA Inference

Once the posterior probabilities are computed, we compute the posterior expectations of two key latent variables, which represent the main output from the LDA algorithm:

1. Posterior vocabulary distribution for each topic: $\{\hat{\beta}_1, \ldots, \hat{\beta}_N\}$

2. Posterior topic mixture for each paragraph in our collection: $\{\hat{\theta}_1, \ldots, \hat{\theta}_D\}$

The first set of output from our LDA procedure identifies the topics. For each topic k, $\hat{\beta}_k = [\hat{\beta}_{k,1}, \ldots, \hat{\beta}_{k,V}]'$, and each entry $\hat{\beta}_{k,j}$ represents the *probability that the word j characterizes topic k*. Our FOMC document collection has $V = 61,432$ unique terms. As a result, each $\hat{\beta}_k$ contains 61,432 entries, the majority of which receives a weight close to zero. Table 1 reports the top 20 words for each topic. This table demonstrates that the topics are clearly identified by the LDA, as the top words from each classified topic are mostly distinct and identify their respective topics with little ambiguity. For example, the first topic consists of keywords such as *policy*, *stance*, etc., indicating that this topic is about monetary policy, and addresses the plan, performance, and outlook of monetary policies. The second topic consists of keywords such as *inflation*, *energy*, etc., suggesting that this topic is about inflation. In fact, the rest of the topics can be similarly identified by the top keyword from their respective classification, as 3) market, 4) employment, 5) economic growth, 6) foreign trade, 7) consumption, and 8) production and investment.[8]

The second set of output is the collection of paragraph-level topic mixture vectors, $\{\hat{\theta}_1, \ldots, \hat{\theta}_D\}$. From this collection, each paragraph d has one mixture, $\hat{\theta}_d = [\hat{\theta}_{d,1}, \ldots, \hat{\theta}_{d,N}]'$. Because there are 8 topics, each vector $\hat{\theta}_d$ has 8 entries, where each $\hat{\theta}_{d,n}$ corresponds to the *proportion* of paragraph d that is devoted to topic n. The 8 entries sum up to one for each paragraph. We plot the time series of the proportion of each topic in Figure 2.1. The shaded areas in Figure 2.1 corresponds to NBER-designated recession periods.

Interestingly, Figure 2.1 shows significant time variation in the proportion of the FOMC minutes devoted to each topic. For example, from 1992 onward, a progressively smaller proportion of the minutes has been devoted to the *growth* topic, which went from the predominant topic in the minutes to a much less prominent portion. At the same time, this decrease has been offset by increases in the proportions of the other topics, particularly those on *policy*, *inflation* and *market*. This finding likely reflects the dynamic roles and

[8]Because topics 4 to 8 are individual components of economic growth, for ease of interpretation by human readers in some of our cross-validations, and as an additional robustness check, we group them into one *economic growth* topic. This results in 4 major topics: *policy*, *inflation*, *growth*, and *market*. The results for tests using this grouping can be found in the Appendix.

responsibilities of the Fed over time: on one hand, it has been increasingly transparent and forthcoming about its policy outlook and expectations. On the other hand, it is increasingly taking up a regulatory role in maintaining the stability of the financial markets, such as negotiating the rescue of systematically important banks and the subsequent TARP initiatives (the proportion of the *market* topic tripled during the recent financial crisis). Overall, this table demonstrates that the FOMC minutes are not uniformly-written documents that always address one particular issue, but a compendium of discussions on various issues, whose proportion change continuously over time. This highlights the importance and usefulness of our topic-based approach.

Finally, as an additional robustness check, we randomly select 50 paragraphs from each of the two groups that satisfy the following properties:

1. Paragraphs classified as containing $\geq 99\%$ of a single topic.

2. Paragraphs classified as containing a mixture of two, three and four topics (each topic having a proportion of at least $\geq 10\%$).[9]

A selection of the texts are presented verbatim in Appendix A.2. We then ask a team of 10 human readers, mostly undergraduate students at the University of Pennsylvania, to identify the topic mixtures of these 100 paragraphs, without revealing the LDA classification result. For paragraghs that are identified by the LDA as containing only a single topic, human readers and the LDA agree in 49 of the 50 paragraphs (e.g. they both identify a paragraph into the *policy* topic). For multiple-topic paragraphs, human readers agree with the LDA in 46 of the 50 paragraphs about the number and type of the topic. However, they often have difficulties pinning down the exact proportions of each topic, especially when the number of topics is higher than three. For example, many readers identify the last paragraph (4 topics) of Appendix A.2 as containing 25% of each topic, whereas the LDA offers a more granular topic proportion mix that is potentially more accurate. This test demonstrates two advantages of our automated topic classification approach. First, it offers an accurate

[9]This is done according to our grouping procedure discussed in the previous footnote.

topic classification that is consistent with common intuition. Second, for each document, it offers a granular, time-varying topic mixture that is more accurate than manual reading, thereby potentially minimizing researcher-induced bias.

2.3.3. Extraction of Contents

Having obtained the estimate of topic proportions, we now proceed to extract the contents for each paragraph-topic combination, using a bag-of-words approach similar to Tetlock (2007) and Jegadeesh and Wu (2013). Specifically, for each paragraph, we simultaneously extract the tone and uncertainty of each topic by tabulating the frequency of keywords in the respective tone and uncertainty lexicons. The Tone content consists of the frequency of positive words (*positive tone*), negative words (*negative tone*), and the difference in frequency between positive and negative tonal words (*net tone*) in a comprehensive tonal lexicon that merges the Harvard IV-4 Psychosociological Dictionary[10] and the financial tonal lists developed by Loughran and McDonald (2011). The *uncertainty* content is the frequency of keywords in the "uncertain words" lexicon developed by Loughran and McDonald (2011).[11]

Since each paragraph is a mixture of 8 topics, the topic contents can be summarized in 8 paragraph-level content *Scores*. Specifically, for meeting t, paragraph d, topics $n = 1, \ldots, 8$, and content $c \in \{\text{positive tone}, \text{negative tone}, \text{net tone}, \text{uncertainty}\}$:

$$Score^t_{d,n,c} = \hat{\theta}^t_{d,n} F^t_{d,c}$$

where $\hat{\theta}^t_{d,n}$ is the topic-n proportion estimate for paragraph d from LDA, and $F^t_{d,c}$ is the number of occurrences of content words from the respective tonal or uncertainty lexicons in paragraph d. In addition, we isolate a list of tonal words that are associated with quantity increases and decreases.[12] We reverse the connotation for these words when they are used in the *inflation* topic. For example, *gain* is considered as a positive word by both lexicons.

[10] Available at http://www.wjh.harvard.edu/~inquirer/homecat.htm.
[11] Available at http://www3.nd.edu/~mcdonald/Word_Lists.html.
[12] Available at http://fnce.wharton.upenn.edu/profile/1661/.

However, it should be treated as a negative word when used in the *inflation* context, because a *gain* in inflation is considered negative by the Fed and thus increases the likelihood of tightening actions. In general, a higher Tone Score indicates a more positive/easing or less negative/tightening tone, and a higher uncertainty Score reflects a higher degree of uncertainty in the paragraph.

Next, similar to Jegadeesh and Wu (2013), we aggregate the Scores to the document level as the sum of individual paragraph scores, weighted by the inverse of paragraph length in number of words:

$$Score^t_{n,c} = \sum_{d=1}^{D_t} Score^t_{d,n,c} \left(\frac{1}{T^t_d} \right), \tag{2.6}$$

where T^t_d is the total number of words in paragraph d, and D_t is the total number of paragraphs in Document t. The term $\left(\frac{1}{T^t_d} \right)$ reflects the intuitive notion that the strength of the topic tone is negatively related to overall paragraph length. Longer paragraphs are more difficult to read and process, and are therefore downweighted.

Figure 2.2 plots the 10-period moving average of the document-level Net Tone and Uncertainty Scores of each of the 8 topics over time. For ease of comparison, the Scores are standardized to a mean of zero and standard deviation of one. Similar to the proportions in Figure 2.1, this figure demonstrate the large difference between the topic *contents* over time. Specifically, Panels 2.2(a) and 2.2(b) display the Tone Scores for topics 1-4 and 5-8, respectively. While most Tone Scores seem to be procyclical, some are much more so than the others. For example, most economic growth related topic Scores are procyclical, becoming more positive during boom periods and turning sharply negative during recressions. This is probably not surprising because FOMC members' discussions on this topic is likely based on their review and outlook of the underlying economic conditions, which are likely to be quite persistent. On the other hand, the Score for the *policy* topic seems to lead economic cycles, as it usually turns half way into the cycle before other series changes direction. This suggests that, particularly during bad economic times, the Fed seems to convey its future (easing) policy direction via more positive policy-related languages before actually taking the easing actions. This finding further suggests that the discussions on some topics are probably more informative than others. Our topic-based approach therefore can highlight the informative topics and construct content indices that are deemed important by the market, and therefore useful for predicting future economic conditions and policies. The rest of our paper is focused on assessing this ability.

Panels 2.2(c) and 2.2(d) plots the Uncertainty Scores for the corresponding topics. These figures demonstrate a similarly high degree of variance in the Fed's use of uncertain languages, both over time and across topics. A case in point is during and after the recent financial crisis: the uncertainty level for the financial market topic spikes during the crisis and the ensuing economic recession, while most growth-related topics have seen uncertainty levels peak just as the economy was coming out of the trough. As conditions get better from 2011 onwards, uncertainty for these topics declined. Precisely around the same period, as the Fed prepares to exit the quantitative easing programs, uncertainty level for policy has spiked. This is consistent with the large volume of media reports that although the Fed is more confident about the economy, it is exceedingly cautious about the pace of future tightening. These observation suggests that the Uncertainty Scores can be used in conjunction of the Tone Scores to create powerful predictive indices. In this paper we focus on the market reaction to the Tone Scores, and leave the Uncertainty Scores and its associated predictive analysis to a companion paper.

2.4. Empirical Tests and Results

This section discusses our empirical tests and reviews the test results. We first assess the informativeness of the minutes as a whole, and then measure the relative informativeness of each individual topic and of its content using market reaction. We then relate our topic-specific content Scores to macroeconomic variables and explore the determinants of the scores. In addition, because a short textual statement is also released immediately after each meeting, we compare the informativeness of these statements with the minutes. Finally, we discuss some possible economic mechanism behind our results and explore the source of informativeness from the minutes by analyzing whether the price reactions to the minutes are permanent.

2.4.1. Market Data

We examine the relation between the content of FOMC minutes and changes in aggregate stock market and interest rates to assess the information content of FOMC minutes. We use transaction prices of SPY to measure intraday market returns.[13] We obtain SPY transaction price data from NYSE's intraday Trade and Quote (TAQ) database. We use 3-month LIBOR rate implied by the nearest maturity Eurodollar futures contract as the interest rate measure.[14] We obtain transaction prices of Eurodollar futures from Chicago Mercantile Exchange.

2.4.2. Informativeness of FOMC Minutes As A whole

We first examine whether the Minutes move the market i.e. whether market volatility increases following its release. Specifically, we examine whether the volatility during the release window is larger on release days than on non-release days using the following regression specifications:

$$V_t = a + bL_t + e_t \tag{2.7a}$$

$$V_t = \alpha + \beta L_t + \sum_{k=1}^{20} \gamma_k V_{t-k} + \epsilon_t \tag{2.7b}$$

where V_t is the 15-minute event window market volatility (on both release and non-release days) computed per Equation (2.1c). Within this short window, confounding effects from other macroeconomic variables are negligible, as no other important economic data is likely announced during this window.[15] L_t is a dummy variable that equals to one if a Minutes is released at date t. Each regression uses between 4,343 and 4,363 days of observation.

[13] SPY is a an actively traded ETF that tracks the S&P 500 index.

[14] Implied 3-month LIBOR=100-Eurodollar futures price.

[15] We confirm this by referring to the Bloomberg Economic Calendar of important economic indicator announcements and find no other significant announcements during this window.

In addition, we use daily volatilities, V_{t-k}, $k = 1, \ldots, K$, in the days prior to the release day to address concerns about potentially biased interest rate expectations and control for any mean reversion induced by the minutes' release. Specifically, suppose a Minutes is completely uninformative, but the market expectation about the content of the minutes can be erroneously distorted between the meeting date and the release date, e.g. by interim speeches from other Fed officials. Therefore, when the Minutes is released, the market corrects its wrong expectation, thereby registering a higher than normal volatility. This produces a positive bias on the coefficient estimates of b in Regression (2.7a). However, because interim changes in expectation are also associated with changes in market volatility, we can use the daily volatility of the k-days between the meeting date and the release date to control for the effect of changing expectations. The estimate for β therefore measures the true level of informativeness of the minutes, conditional on all prior expectations.

We fit Regression (2.7b) above using K=0, 5, 10, and 20 trading days. The Minutes are released at 2:00 pm and hence we use the 15-minute window from 2:00 pm to 2:15 pm as the event window. To facilitate interpretation, we scale all regression coefficients by the unconditional mean of V_t across all observations. The coefficient estimates $\hat{b}$ and $\hat{\beta}$ can thus be interpreted as the *incremental* volatility introduced by the release of the minutes as a percentage of the average volatility in the event window across both release and non-release days.

Table 2.3 reports the coefficient estimates. The estimate for the release dummy, L_t, is significantly positive for all specifications, and it ranges from .5919 to .6130 for SPY. These estimates indicate that the volatility on the when the minutes are released is about 60% bigger than that during the same time on other days. The inclusion of lagged volatility as control variables increases regression R^2 since it accounts for time-variation in volatility, but it in does not materially affect the slope coefficients.

Table 2.3 also reports the results for volatility of LIBOR. The slope coefficients for LIBOR are between .2283 and .3054. Therefore, the proportional increase in LIBOR volatility is about half that for SPY, indicating that the minutes have a relatively larger impact on the stock market.[16]

The result is surprising. Both Fleming and Piazzesi (2005) and Lucca and Moench (2015) find that on the actual FOMC meeting days, market volatilities are significantly higher. By contrast, the minutes are released several weeks after the original meetings, and intuitively, every action discussed by the minutes should already be public knowledge by then. Our finding that market volatility is also significantly higher on release days thus indicates that the minutes' *language* itself does contain additional, "soft" information not incorporated in the quantitative policy announcements such as interest rates, nor are they sufficiently conveyed by other post-meeting communications such as speeches and interviews of Fed officials. Recent macroeconomics literature such as Sims and Zha (2006) has used structural models to identify policy changes from observed interest rate data. Our findings indicate that the Minutes contain information beyond the rates data and therefore can be utilized to enhance monetary policy models without additional filtering. We examine the possible economic mechanisms in more detail in Section 2.4.6.

Next, we examine the informativeness of the overall document without dividing it into topics. This analysis sets a benchmark to judge the incremental information that can be divined through topic level analysis. We compute the Tone Scores for the entire document level relate them to unexpected market volatilities in the following regression:

$$UV_t = \alpha + \beta_c Score_c^t + \epsilon_t \tag{2.8}$$

where UV_t is the unexpected volatility around the event window on release date t, computed per Equation (2.2). $c \in \{\text{net tone}, \text{positive tone}, \text{negative tone}\}$ are the document-level net, positive, and negative tone Scores, computed per Equation (2.6) while setting all $\hat{\theta}$'s

[16] Lucca and Moench (2015) demonstrate large excess returns in equities in anticipation of ammouncements after FOMC meetings. As an additional robustness check, we also extend the release window to 20 minutes from 1:55pm to 2:15pm to account for any pre-release leakage of information, or anticipation of such information. The results within the 20-minute window, shown in Panel B of the same table, are similar to that within the 15-minute window.

equal to one. Each regression uses 138 observations corresponding to the minutes' release dates. These regressions explore the relation between the overall document tone and market reaction. If, for example, a more positive overall tone is more informative, then we would expect a positive correlation between tone and volatility, i.e. a positive estimate for β_{net}.

Table 2.4 reports the regression estimates. None of the coefficient estimates for the document-level content Scores are statistically different from zero, with t-statistics ranging between -1.69 and 0.32. This suggests that, on the document level, the market does not perceive the tone Scores as useful, as neither more positive nor more negative tone Scores are associated with higher market volatility. This suggests that when the entire document is viewed as a single unit the document tone is not related to changes in volatility. It is, however, possible that some topics are informative than others and the informative topics may not be evident when all topics are simultaneously considered. Therefore, our next set fore set of tests examine the informativeness of individual topics.

2.4.3. Relative Informativeness of Individual FOMC Topics

Although the tone of the entire document is not informative, it is possible that some of the individual topics may be informative while some are not. For example, our discussions with industry practitioners reveal that they consistently find the discussion on inflation to be more informative than that on economic growth. Our next set of tests evaluates the informativeness of individual topics.

We examine the relation between unexpected volatility the proportion of each topic, and we also assess the informativeness of each topic's contents using Tone Scores. We specify the following relations between topic proportion, content Scores, and event window return volatility for each $c \in \{\text{positive tone}, \text{negative tone}, \text{net tone}\}$:

$$UV_t = a + \sum_{n=1}^{8} b_n \hat{\theta}_{n,t} + rX_t + e_t \tag{2.9a}$$

$$UV_t = \alpha + \sum_{n=1}^{8} \beta_n Score_{c,n}^t + \gamma X_t + \epsilon_t \tag{2.9b}$$

where UV_t is the unexpected volatility around the event window on release date t, computed per Equation (2.2), and X_t is the vector of macro controls variables that include:

- *IntRate*: the latest daily closing yield of 10-year Treasury notes obtained from the Department of Treasury.
- *UnEmp*: latest monthly rate of unemployment obtained from the Bureau of Labor Statistics.
- *Recession*: a dummy variable which equal to one if meeting date t falls within a NBER-designated recession period.

In order to explore market reactions as broadly as possible, we fit each regression with volatility data computed from both equity (SPY) and debt/interest rate (Eurodollar) markets. Each regression uses 138 daily observations from 2000 to 2015. In this setting, an estimate of b_n or $\beta_{c,n}$ that is statistically different from zero indicates informativeness of a topic, or its content Score: a significantly positive $\hat{\beta}$ suggests that the market respond more to a more positive topic tone while a significantly negative $\hat{\beta}$ suggests that the market find more information in a more negative topic tone. Similarly, a significantly positive $\hat{b}$ for topic n indicates that the market finds the discussion of this topic informative, when it is discussed more, regardless of the tone.

Table 2.5 displays the coefficient estimates from the proportion Regression (2.9a). The *growth* topic is omitted from the regression to prevent multicolinearity. All independent variables in the regressions are standardized to mean zero and unit standard deviation. First, relative to the *growth* topic, the coefficient estimates for *policy*, *inflation*, and *employment* proportions are all statistically significant and positive. This is consistent with the findings in Table 2.3 that the minutes do contain additional information. The findings in this table identify the granular source of this information: the market focuses it attention on the languages on monetary policy, inflation, and employment situations, and do not pay particular attention to growth-related discussions. As the discussion on these key topics becomes more detailed (thus higher proportions), more information is transcribed in the texts, and the market responds more.

Columns (5) through (8) present the slope coefficients of Regression (2.9a) where the dependent variable is directional change in SPY or LIBOR. The slope coefficient for SPY is significantly positive and for LIBOR is significantly negative for tone on policy. A larger proportion of policy oriented discussion seems to be correlated with the Fed easing interest rates, which in turn results in lower rates and higher stock market. The proportion of other topics are not related to directional changes in SPY or LIBOR.

The left four columns of Tables 2.6 to 2.8 report the coefficient estimates for the tone Scores in Regression (2.9b). These tables suggest that, in addition to topic proportions, the market also views the tones of different topics differently and assigns different informational value to them. First, the coefficient estimates for inflation topic's Net tone Score is significantly negative, indicating a more negative or less positive tone is associated with a higher magnitude of market reaction. This is further confirmed by the positive estimate in Table 2.7, which shows that more negative language in inflation discussions is indeed associated with higher unexpected volatilities. Moreover, the estimate for positive inflation tone is not significant, further suggesting that market participants are particularly looking out for negative discussions on inflation. A similar pattern can be found for the policy and unemployment topics. For the policy topic in particular, the estimates using Eurodollar volatilities are more significant than using SPY volatilities. This indicates that the short-term debt markets are more sensitive to the discussions on monetary policy than the equity market. Broadly speaking, these results are consistent with the notion that the market reacts stronger to unanticipated tightening actions (indicated by more negative discussions) than an easing policy stance.

Our next set of tests explore the directional impact of our topic Scores. Here we examine whether, for example, a more negative discussion on inflation moves the market up or down. After all, such discussion could indicate bad current conditions, but at the same time signal future easing actions. If the market is forward looking, then its response would not be uniform. As such, the informativeness of the topics is reflected by not only by market volatility, but also from the relation between the Scores and raw, directional returns during the event window. We therefore fit the following regression:

$$R_t = \alpha + \sum_{n=1}^{8} \beta_n Score^t_{\text{net tone},n} + \gamma X_t + \epsilon_t, \tag{2.10}$$

where R_t is the 15-minute event-window equity and interest rate market returns constructed according to Equations (2.1a) and (2.1b) and X_t is the vector of controls used in the previous regression. This regression explores the micro relation between the topic tones and directional returns. If the market indeed thinks that a particular tone for a particular topic is good/bad news, then it should respond accordingly, resulting in a positive/negative estimate for β.

The right four columns of Tables 2.5 and 2.6 report the coefficient estimates for the topic proportion and Net Tone Scores, respectively. First, surprisingly from Table 2.6, the estimates for the financial market topic is significantly negative: The equity market in particular actually interprets a more positive tone of market discussion as bad news, assigning a 0.1% lower with a one-standard-deviation change in the tone. This suggests that perhaps a need to prop up the economy is more positive in tone but the market views it as a negative signal. More significantly, the estimates for the policy and inflation topics are significantly positive for both SPY and Eurodollar markets (for the Eurodollar market, a negative coefficient indicates positive price movement): market return is on average between 0.09% and 0.13% higher during the 15-minute window with a one-standard-deviation increase in the scores. Thus, for discussions on monetary policy and inflation, the markets do view more positive tones as good news.

Finally, note that separately using equity and debt market data in directional regressions allows us to interpret the exact meanings of "positive" and "negative" in FOMC languages. The logic is as follows: while tightening actions might not have as pronounced an impact on equity markets, they impact the credit markets more directly, because increases in interest rates (or rate expectations) is directly translated to higher yields. From Table 2.6, the coefficients for SPY and Eurodollars are indeed opposite in most cases: a more negative policy tone, for example, is associated with positive yield changes and negative stock market returns. This finding suggests that our topic-content Scores capture the degree of policy "tightness": a more positive tone is interpreted as a move toward easing, while a more negative tone means policy tightening.

2.4.4. Determinants of Topic Proportion and Tone

How "orthogonal" are our granular, text-based measures from existing economic indicators? After all, the Fed is likely to take into account current economic conditions when formulating monetary policies. In addition, a whole section of the Minutes is devoted to reviewing current economic conditions and providing an outlook for future conditions. Many theoretical and structural frameworks of monetary policy making, for example the Taylor rule, also stiputates that monetary policy, usually proxied by the nominal interest rate, is related to changes in economic variables such as output, inflation and unemployment. Does the Taylor rule matter when the FOMC members are in the meeting room? This subsection specifically examines the relation between the proportion and content Scores of each topic and current economic conditions. From our discussion in Section 2.3.3, we expect the tone of several growth-related topics to be procyclical and follow the traditional Taylor rule, while some other, more "forward-looking" topics might not be the case. For example, as the FOMC members have much latitude in their policy discussions, the effect of macro variables on the Policy Score is likely to be ambiguous: if the Fed correctly anticipates economic cycles and changes policy before the cycle changes, then we might not see a significant relation between policy proportion/content and contemporaneous macro variables. The Policy Score is therefore likely to be the most orthogonal among the topics.

We examine the determinants of the topic proportion and content Scores via the following regressions:

$$\hat{\theta}_{n,t} = a + bIntRate_t + rUnEmp_t + dRecession_t + e_t \tag{2.11a}$$

$$Score^t_n = \alpha + \beta IntRate_t + \gamma UnEmp_t + \delta Recession_t + \epsilon_t \tag{2.11b}$$

where $Score^t_n$ is the Net Tone Content Score for topic n and Minutes t, computed per Eq. (2.6). $\hat{\theta}_{n,t}$ is the topic-n proportion in Minutes t estimated using Eq. (2.5). We fit the regression using all 176 minutes Documents from 1991 to 2015.

Table 2.9 presents the coefficient estimates for the proportion regressions. The proportion of most growth-related topics are positively related to interest rates and negatively related to unemployment. The opposite relations can be found for the *inflation* and *market* topics. This suggests that during bad times, the Fed is more concerned about inflation (or deflation) and the health of the financial markets, than for economic growth and sub-topics like foreign trade. Table 2.10 reports the coefficient estimates for the Tone Score regressions. And here again, the tone of most *growth*-related topics are procyclical. Interestingly, the tone of the *policy* topic is not significantly related to existing economic conditions: none of the coefficient estimates are statistically significant. This again highlights the fact that the policy discussions during the meetings probably incorporate factors beyond current economic conditions, and therefore, the *policy* topic can serve as a leading indicator of the economy, which is corroborated by Fig. 2.1, where its Score usually "flips" half way into the economic cycle. We explore the predictive power of the topic Scores in a companion paper.

2.4.5. Relative Informativeness of Statements vs. Minutes

Another useful test in illustrating the efficacy of our granular information extraction approach is comparing the minutes-based Scores with the information contained in the languages of meeting *statements*, which are very short documents (usually one paragraph) released immediately after each meeting. These statements outline the policy decision made during the meeting and (for later years) very succinctly discuss the rationale for such decisions. As such, the languages of the statements can also potentially contain incremental information not conveyed by the hard numbers. Because of their very short length, the statements are not suitable for topic-based analyses. We therefore compute the content Scores for these statements as a whole, then compare those with our more granular Scores from the minutes and examine whether the granular Scores contain yet another layer of incremental information in addition to those contained in the statements.

Panel A of Table 2.14 assesses the incremental informativeness of statement languages using market reactions on the day of the meeting, and on the day of the corresponding minutes' release. Not surprisingly, the tone of the statements is significantly related to market reaction on the meeting day, even after controlling for any interest rate changes made during the meeting. This is not the case on the minutes' release days, as the statement tone is not statistically significant for either raw or unexpected volatility regressions.

To see this more clearly, we relate the informativeness of the statement languages to that of the minutes' individual topics in a predictive setting. If statements are as informative as the granular minutes-based Scores, then their tone should be able to predict the topic tone scores from the corresponding Minutes released for the same meeting. We therefore modify Regression (2.11b) as follows:

$$Score_n^t = \alpha + bScore_{\text{statement}} + \beta IntRate_t + \gamma UnEmp_t + \delta Recession_t + \epsilon_t, \qquad (2.12)$$

where $Score_{\text{statement}}$ is the overall Net Tone Score for the statement released at the same meeting.

Panel B of Table 2.14 presents the coefficient estimates. With the exception of the inflation topic, the coefficients for statement Tone Score are not statistically significant in all topics. This indicates that, although informative on their own, the statement tones are not enough to predict the tone of individual minutes topic Scores, and the more granular scores contain information not captured by the languages of the statements.

2.4.6. Discussion: Is This Real Information?

This section explores the reason why the minutes are informative. On the surface, this seems puzzling: the minutes are released a long time after the meetings, why would they contain any incremental information at all? However, there is an important distinction: the staff economists at the Fed and the FOMC members have access to a wide variety of confidential economic data, such as detailed records of interbank lending, that are not observable to other researchers. It is likely, therefore, that their information set is superior to that of other market participants. Due to the confidential nature of the data, they cannot disclose any quantitative facts in the minutes. However, it possible that such "inside" information influences the tone and uncertainty level of the minutes' language.

We can jointly test the above hypothesis and whether such "soft" information is transmitted to the market by observing the market reactions to the minutes: if there is new information about the future economy *and* the information is transmitted to the market through the minutes, then the effect on prices should be permanent rather than temporary, and the price "jump" on the minutes release day that we document in the last section should be persistent and not revert quickly. In other words, because temporary price changes would be followed by price changes in the other direction, if market volatility declines after the minutes' release windows, this would suggest that some real information that would have flowed to the market after the release windows indeed is revealed to the market during the release windows.

Figure 2.3 plots the average minute-by-minute return volatility in 15-minute bins for both release and non-release days. First, we confirm the same pattern found in the treasury market: market volatility spikes dramatically to about 1.6 times the normal levels on days where FOMC minutes are released. More importantly, volatility quickly declines after the minutes are released to about 20% lower than non-release days. As a result, the initial price jump at the release do not on average revert back, and prices on average stay at the new levels. This permanent "shift" in prices indicates that the overall level of uncertainty in the market is lower after the release of the minutes, and supports our hypothesis that information is indeed transmitted from the Fed to the markets through the text of the minutes, and in a permanent fashion.

2.5. Alternative Specifications and Robustness Checks

Because our granular, topic-based textual content scores are derived using new methodology, we conduct a series of robustness checks to ensure that the economic inferences that we, and future researchers, can draw from our methods are valid and broadly applicable, and are not subject to the variations in test specifications and peculiarities of the text samples. This section outlines some concerns that one might raise about using our LDA approach to classify the FOMC minutes, and the results of our additional tests to address these concerns.

2.5.1. Shifts in Textual Sample Over Time

One might worry that over time, the writing style of the minutes might dramatically change, thus making our approach more prone to capturing style changes rather than variations in actual information. One particular example from our discussion with Fed personnel is that, after 2011, the minutes became much longer and more detailed in many topics. In addition, around the same time, the Fed began to release the actual economic forecasts by individual FOMC members at 4 out of the 8 meetings every year. This setting allows us to test the robustness of our methodology across different periods with different writing styles and potentially different overall informativeness. We first separate our sample into two halves, before and after (including) 2011. We then examine whether the overall informativeness has changed by separately plotting the average volatility around release days, for both samples, in the top panel of Figure 2.3.

This graph demonstrate that post-2011, the market reaction to the minutes' release is stronger, with the average volatility about 200% higher than normal. While this suggests that the market does pay increased attention to the minutes, Table 2.13, replicating Regression (2.9b) for the two subsamples, shows no change in the *relative* informativeness of individual minutes topics captured by our granular Tone Scores: both the magnitude and the statistical significance of the estimates are similar across subsamples. This result suggests that our methodology is stable even when the overall level of informativeness can change with the writing style of the minutes.

In addition, we further separate the post-2011 sample into two subsample of release days according to whether the preceding meetings are accompanied by the release of Summary of Economic Projection (SEP) materials. After the April 2011 meeting, the FOMC begins to release participants' three-year and long-run projections of three economic indicators and target fed funds rate, based on their "individual assessments of appropriate monetary policy". These projections are released immediately after four of eight meetings annually. If the language of the minutes contain similar information to the projections, then for meetings with SEP releases, the market reaction to the subsequent release of the minutes would be more muted. The bottom panel of Figure 2.3 plots the daily volatility levels on release days with and without SEP releases. The figure shows that volatility levels are similar on both types of days, indicating that releasing individual forecasts does not decrease the relative informativeness of the minutes, and that the information contained in the minutes' languages is deemed by the market to be orthogonal to the SEP data.

2.5.2. Alternative Lexicons and Tone Measures

One might also worry that, as our topic classification becomes more granular, the results are more sensitive to small changes in the tone measures that are purely attributable to the construction of the tonal scores. To address this issue, we conduct two tests where we intentionally magnify and reduce such differences. First, we modify Regressions (2.9b) and (2.10), using the change of the topic Net Tone Scores, rather than the Scores themselves, as the independent variable:

$$UV_t = \alpha + \sum_{n=1}^{8} \beta_n \left(Score_{c,n}^{t} - Score_{c,n}^{t-1}\right) + \gamma X_t + \epsilon_t \tag{2.13a}$$

$$R_t = \alpha + \sum_{n=1}^{8} \beta_n \left(Score_{c,n}^{t} - Score_{c,n}^{t-1} \right) + \gamma X_t + \epsilon_t. \tag{2.13b}$$

This setting potentially introduces more noise into the analysis, as the differences in tone *between* meetings can be affected by both actual information and mechanically by the construction of our Scores. Beside serving as a robustness check, this specification also serves to examine whether the *change* in tone is related to volatility in the markets. Similarly, we artificially dampen such differences by inserting the absolute value of tones as the regressors in the above specification:

$$UV_t = \alpha + \sum_{n=1}^{8} \beta_n \left| Score_{c,n}^{t} \right| + \gamma X_t + \epsilon_t \tag{2.14a}$$

$$R_t = \alpha + \sum_{n=1}^{8} \beta_n \left| Score_{c,n}^{t} \right| + \gamma X_t + \epsilon_t. \tag{2.14b}$$

Beside serving as a robustness check, this specification also serves to examine whether the *volatility* in tone is related to volatility in the markets. Table 2.12 presents the coefficient estimates for regressions with tone differences and Table 2.11 presents results using the absolute value of tones. For tone changes, the results are very similar in signs and slightly larger in magnitudes. This is consistent with the intuitive notion that large changes in tones attracts more attention than smaller changes. Similarly, Table 2.11 shows that the coefficients for the volatility of key tone Scores from Table 2.6 are significantly positive, again intuitively confirming that higher variations in tone are indeed associated with more market reaction. Furthermore, the fact that the results are qualitatively unchanged from those in Table 2.6 indicates that our method is not subject to mechanical noises introduced by the construction of tone Scores.

Another concern is that, although we use the Loughran and McDonald (2011) tonal lexicons as part of our main lexicon, there might still be ambiguity in the *interpretation* of the tone of some words classified as positive or negative by the lexicons. To address this issue, we first recompute the Net Tone Score for each topic using a unified, weighted lexicon also used by Jegadeesh and Wu (2013). This dictionary is constructed by merging the Harvard and LM lexicons (both positive and negative words), but instead of assigning any tonal connotations, each word is weighted objectively according to the market reaction to the 10-K filings in which a word is used. In this sense, a word associated with negative market returns is classified as a negative word. We then replicate Regressions (2.9b) and (2.10) using this new Net Tone Score and present the results in Table 2.15. Again, the results are very close to the original specification. This suggests that our approach is robust to alternative choices of lexicons.

2.6. Conclusion

We present a novel approach in financial content analysis to determine the informativeness of FOMC meeting minutes. This automated approach is based on the Latent Dirichlet Allocation (LDA) algorithm, which enables us to dissect minutes into distinct topics and simultaneously extract the tone and uncertainty level of each topic. In an event study setting, we use market reaction to assess the relative informativeness of each topic and find a significant relation between topic contents and market volatility. Furthermore, we find evidence consistent of the Fed possessing superior information, which is transmitted to the market through the text of the minutes.

Our measures of economic and policy outlook/uncertainty are both model-independent and robust, and can be readily applied to structural macroeconomic models, as well as reduced-form predictive models where policy uncertainty serves as an input. We are currently exploring several of such these issues.

Figure 2.1: FOMC Topic Proportions Over Time

This figure plots the proportion of each topic identified by the LDA algorithm for each of the 196 FOMC minutes released between 1990 and 2015. The raw inputs for this figure are the 5,644 paragraph-level topic mixture vectors, $\{\hat{\theta}_1, \ldots, \hat{\theta}_{5644}\}$. Each vector $\hat{\theta}_d$ has 8 entries, where each $\hat{\theta}_{d,n}$ corresponds to the *proportion* of paragraph d that is devoted to topic n. The 8 entries sum up to one for each paragraph. The document-level proportion are paragraph-level proportions weighted by paragraph length. The shaded areas correspond to NBER-designated recession periods.

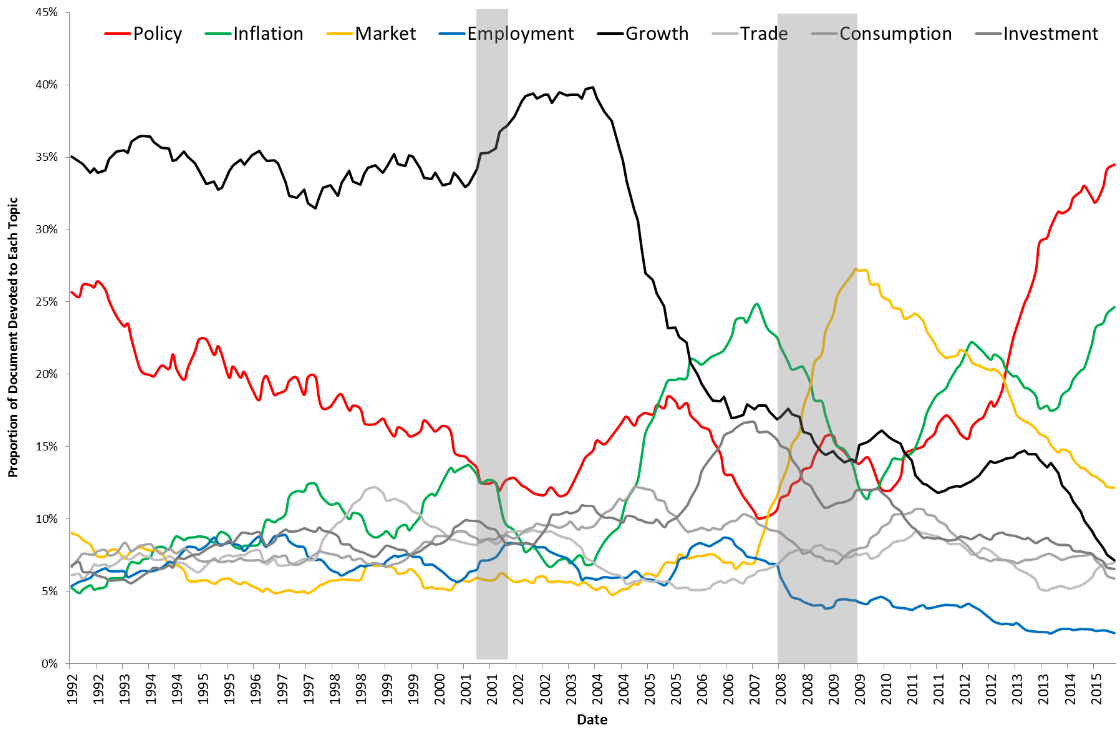

Figure 2.2: FOMC Topic Content Scores Over Time

This figure plots the Net Tone and Uncertainty Scores of each topic identified by the LDA algorithm for each of the 196 FOMC minutes released between 1990 and 2015. The raw inputs for this figure are the 5,644 paragraph-level topic mixture vectors, $\{\hat{\theta}_1, \ldots, \hat{\theta}_{5644}\}$. Each vector $\hat{\theta}_d$ has 8 entries, where each $\hat{\theta}_{d,n}$ corresponds to the *proportion* of paragraph d that is devoted to topic n. These proportions are used to compute document-level Net Tone and Uncertainty Scores according to Equation (2.6) of the text. The shaded areas correspond to NBER-designated recession periods.

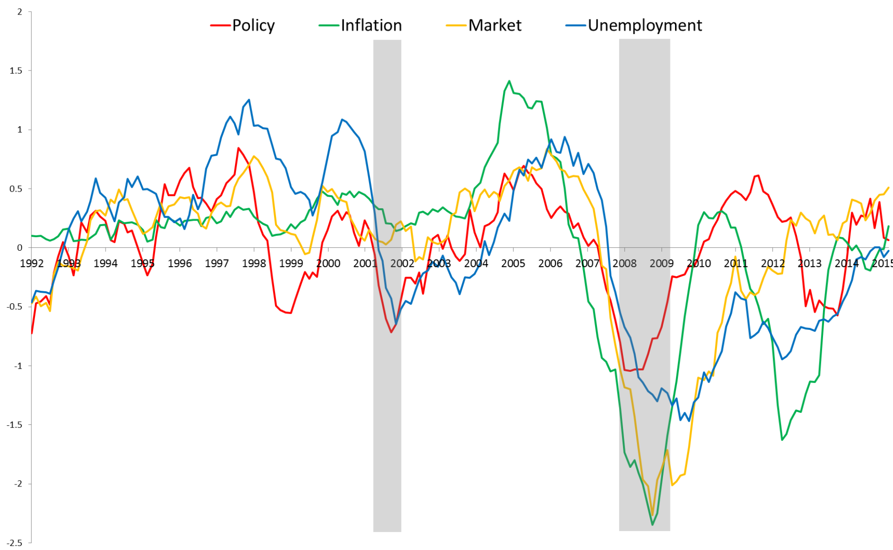

(a) Topic Net Tone Scores Over Time, Topics 1 to 4

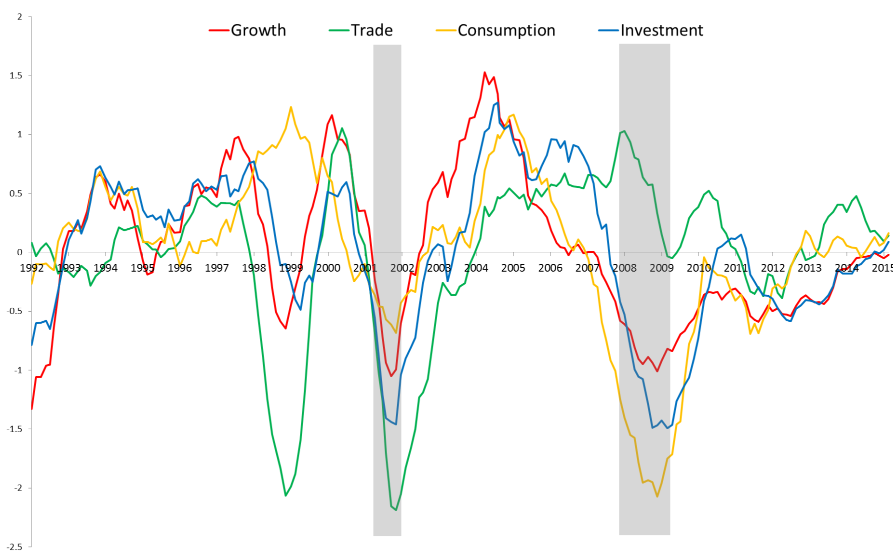

(b) Topic Net Tone Scores Over Time, Topics 5 to 8

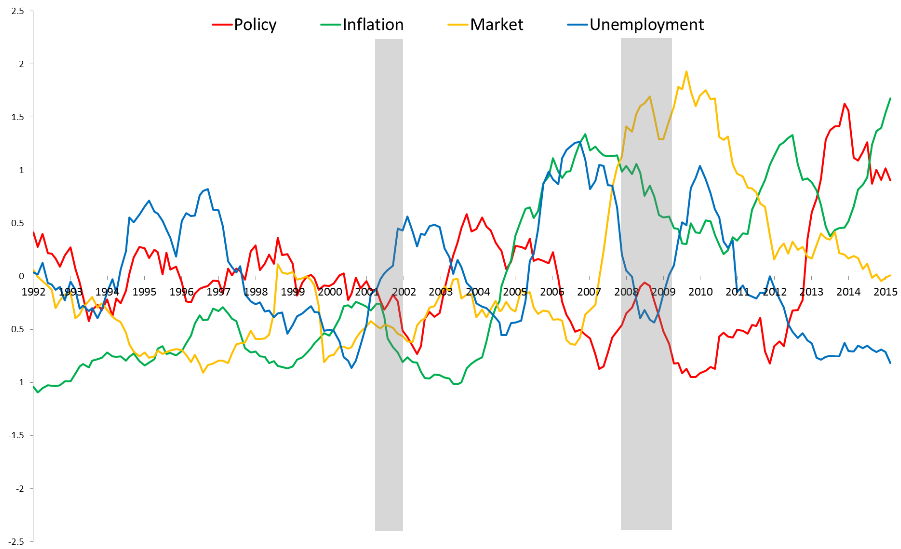

(c) Topic Uncertainty Scores Over Time, Topics 1 to 4

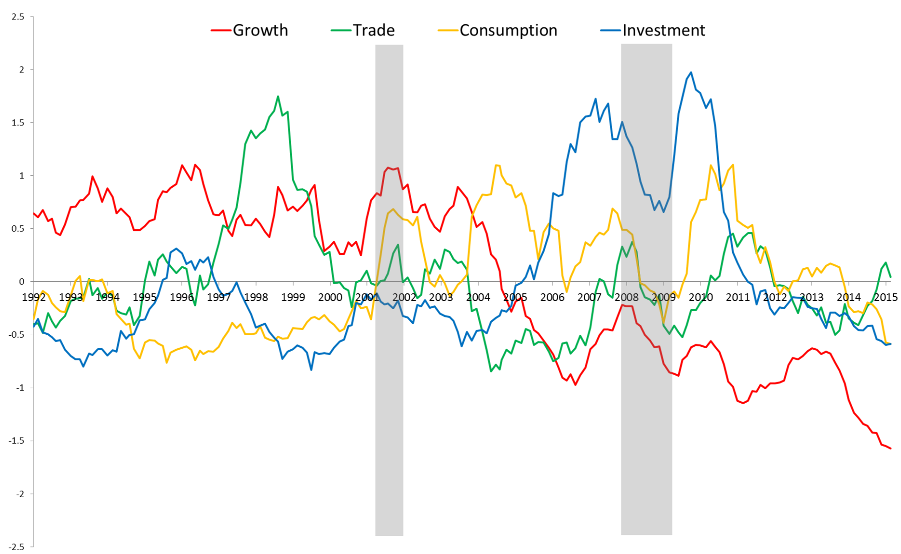

(d) Topic Uncertainty Scores Over Time, Topics 5 to 8

Figure 2.3: Market Reaction to Release of FOMC Minutes

This figure plots the daily average of 15-minute raw SPY return volatility, in various subsamples, from $t-3$ to $t+2$ days around the minutes release days t. The volatilities are computed according to Eq. (2.1c) of the text. The top panel shows the ratio of volatility on release days over that on non release days, for the full sample between 2000 and 2015, and two subsamples of 2000-2011 and 2011-2015, respectively. The bottom panel plots compares the raw volatility levels in the post-2011 subsample, between meetings with and without the release of Summary of Economic Projections (SEP) data.

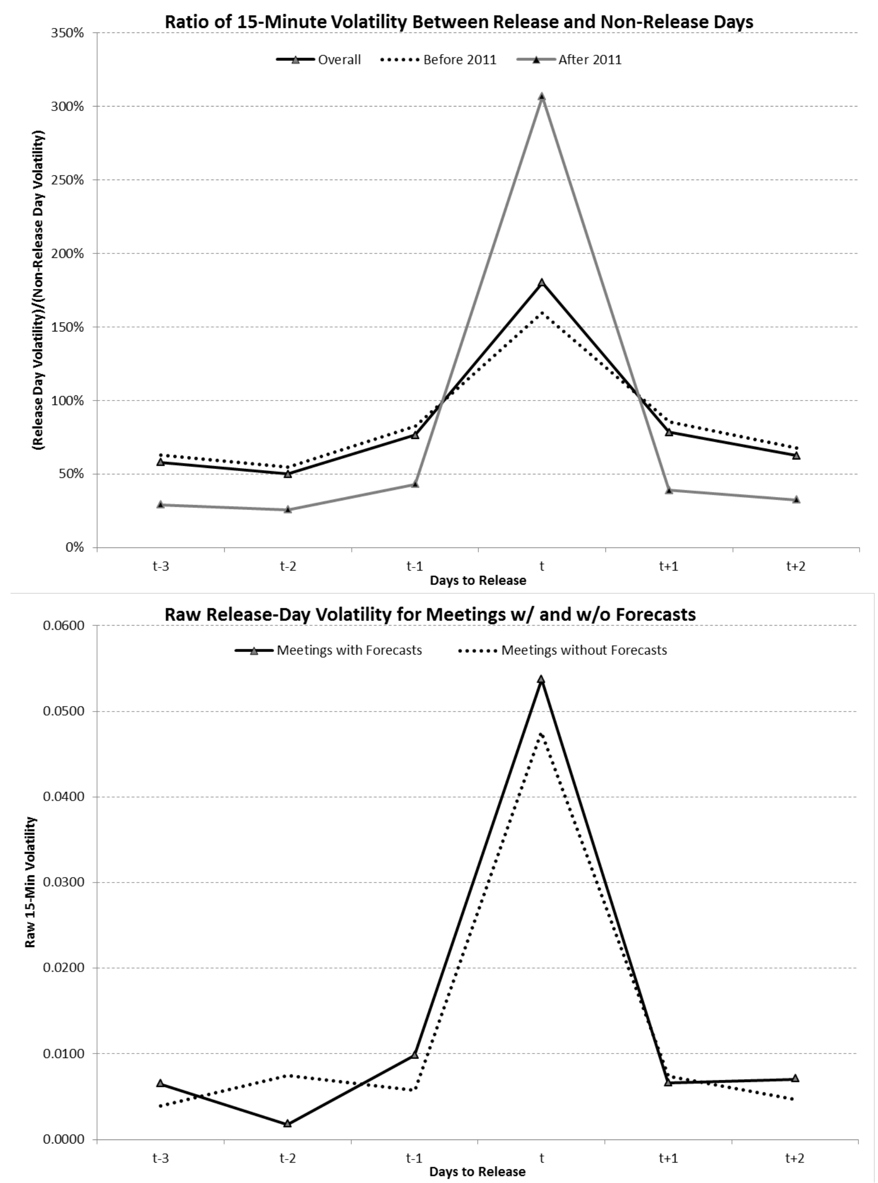

Figure 2.4: Temporary vs. Permanent Reaction to Release of FOMC Minutes

This figure plots the ratio of average return volatility in 15-minute bins between release and non-release days. Return volatility is calculated as the standard deviation of minute-by-minute returns in each 15-minute bins according to Eq. (2.1c) of the text. The sample period is 2000-2015.

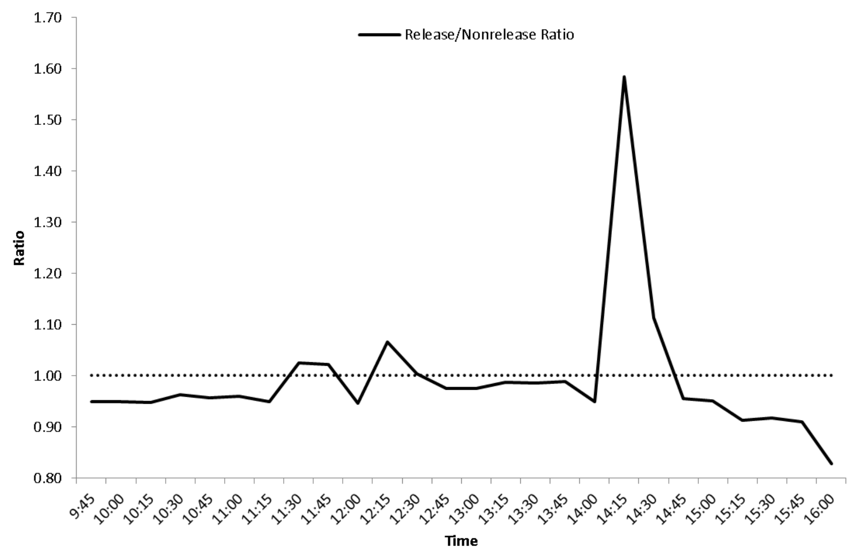

Table 2.1: Distribution of Top LDA Topic Keywords

This table reports the top 20 words for each topic identified by the LDA procedure. Each column in this table represents a topic $k = 1, \ldots, 8$, and the weights are estimates of $\hat{\beta}_{k,j}$ and represent the *probability that the word j characterizes topic k*. The distributional assumptions for the LDA model are outlined in Section 2.3 of the text. The estimation uses 5,644 paragraphs from FOMC meeting minutes released between 1990 and 2015.

Topic 1		Topic 2		Topic 3		Topic 4	
Weight	Word	Weight	Word	Weight	Word	Weight	Word
0.0445	policy	0.0788	inflation	0.0240	market	0.0335	labor
0.0216	monetary	0.0265	energy	0.0206	credit	0.0295	employment
0.0188	funds	0.0255	consumer	0.0172	yields	0.0250	job
0.0143	reserve	0.0226	labor	0.0146	financial	0.0247	workers
0.0133	risks	0.0212	core	0.0144	liquidity	0.0231	payroll
0.0113	financial	0.0178	expectations	0.0142	loans	0.0157	manufacturing
0.0104	agreed	0.0119	compensation	0.0141	securities	0.0151	hiring
0.0100	directive	0.0111	pce	0.0126	debt	0.0147	nonfarm
0.0086	guidance	0.0108	food	0.0123	spreads	0.0138	private
0.0080	purchases	0.0103	unemployment	0.0112	equity	0.0116	unemployment
0.0074	target	0.0099	real	0.0109	corporate	0.0108	inflation
0.0071	stability	0.0090	costs	0.0107	funds	0.0104	hourly
0.0071	easing	0.0089	index	0.0106	commercial	0.0103	services
0.0068	consistent	0.0085	commodity	0.0098	bank	0.0101	earnings
0.0065	stance	0.0082	oil	0.0086	nonfinancial	0.0099	food
0.0063	expectations	0.0072	slack	0.0078	investors	0.0095	costs
0.0057	tightening	0.0069	producer	0.0077	institutions	0.0091	force
0.0056	asset	0.0067	reflecting	0.0075	lending	0.0087	output
0.0054	action	0.0065	subdued	0.0072	issuance	0.0085	utilization
0.0052	view	0.0065	headline	0.0071	bonds	0.0085	construction
Topic 5		**Topic 6**		**Topic 7**		**Topic 8**	
Weight	Word	Weight	Word	Weight	Word	Weight	Word
0.0208	economy	0.0340	foreign	0.0448	consumer	0.0447	production
0.0169	business	0.0315	exports	0.0381	sales	0.0369	manufacturing
0.0129	economic	0.0289	u.s	0.0335	housing	0.0354	inventories
0.0111	demand	0.0268	dollar	0.0168	homes	0.0275	output
0.0087	productivity	0.0223	imports	0.0165	mortgage	0.0266	motor
0.0076	investment	0.0219	economies	0.0164	starts	0.0223	investment
0.0072	pressure	0.0166	countries	0.0145	construction	0.0201	industrial
0.0068	firms	0.0152	trade	0.0138	income	0.0160	sales
0.0063	financial	0.0140	major	0.0135	household	0.0149	equipment
0.0058	fiscal	0.0128	currencies	0.0134	gains	0.0143	vehicles
0.0057	prospects	0.0125	industrial	0.0131	expenditures	0.0141	business
0.0056	capital	0.0118	deficit	0.0105	single-family	0.0136	stocks
0.0055	confidence	0.0117	united	0.0101	retail	0.0118	wholesale
0.0055	strength	0.0112	japan	0.0098	motor	0.0118	capacity
0.0054	sectors	0.0098	exchange	0.0097	personal	0.0112	utilization
0.0053	potential	0.0097	euro	0.0091	purchases	0.0097	ratio
0.0051	favorable	0.0088	emerging	0.0078	vehicles	0.0090	industries
0.0050	costs	0.0084	sovereign	0.0077	existing	0.0087	retail
0.0049	anecdotal	0.0080	abroad	0.0076	residential	0.0074	accumulation
0.0049	stimulus	0.0072	european	0.0073	sentiment	0.0072	factory

Table 2.2: Distribution of Estimated Weights Among Top Topic Keywords

This table reports the sum of weights for the top 10 words, as well as sums of 10-word bins up to word 50, and the sums of words 51-100 and 101-200. Each column in this table represents a topic $k = 1, \ldots, 8$, and the weights are estimates of $\hat{\beta}_{k,j}$ and represent the *probability that the word j characterizes topic k*. The distributional assumptions for the LDA model are outlined in Section 2.3 of the text. The estimation uses 5,644 paragraphs from FOMC meeting minutes released between 1990 and 2015.

Top Words	Sum of Weights							
	(1) Policy	(2) Inflation	(3) Market	(4) Employment	(5) Growth	(6) Trade	(7) Consumption	(8) Investment
Top 10	**16.26%**	**22.74%**	**13.70%**	**17.67%**	**11.00%**	**21.72%**	**20.74%**	**24.86%**
11-20	7.20%	8.66%	8.75%	9.33%	5.90%	10.74%	9.71%	10.09%
21-30	4.95%	5.88%	5.78%	7.33%	4.44%	6.35%	6.32%	6.26%
31-40	3.66%	4.66%	4.60%	5.53%	3.91%	5.12%	4.87%	4.71%
Top 50	**35.34%**	**45.56%**	**36.52%**	**44.31%**	**28.65%**	**48.45%**	**45.64%**	**49.70%**
51-100	12.55%	14.19%	12.57%	15.40%	12.68%	14.99%	12.55%	13.00%
Top 100	**47.89%**	**59.75%**	**49.10%**	**59.71%**	**41.33%**	**63.44%**	**58.19%**	**62.70%**
101-200	14.53%	15.93%	14.50%	14.72%	15.34%	13.90%	14.15%	14.13%

Table 2.3: Market Reaction to the Release of FOMC Minutes

This table reports the coefficient estimates for Regression (2.7b), fitted using the release-day dummy, as well as lagged volatilities using K=0, 5, 10, and 20 trading days. The equity market regression uses transaction prices of SPY to measure intraday market returns. The Eurodollar market regression uses 3-month LIBOR rate implied by the nearest maturity Eurodollar futures contract as the interest rate measure. The top panel reports results using 15-minute event window market volatility and the bottom panel reports results using 20-minute window. Each regression uses between 4,343 and 4,363 days of observation.

Panel A. 15-Minutes Volatility

Equity Market	Number of Lags in Control			
	(None)	(5)	(10)	(20)
Release Dummy	0.5919***	0.6081***	0.6130***	0.6032***
	(6.32)	(6.99)	(7.13)	(7.02)
No. Obs	4363	4358	4353	4343
adj. R-sq	0.011	0.149	0.17	0.182
Eurodollar Market	Number of Lags in Control			
	(None)	(5)	(10)	(20)
Release Dummy	0.3054*	0.2283	0.2665*	0.2638*
	(2.02)	(1.74)	(2.01)	(1.98)
No. Obs	2627	2622	2617	2607
adj. R-sq	0.004	0.124	0.136	0.151

Panel B. 20-Minutes Volatility

Equity Market	Number of Lags in Control			
	(None)	(5)	(10)	(20)
Release Dummy	0.4965***	0.5174***	0.5418***	0.5084***
	(5.17)	(5.77)	(6.14)	(5.76)
No. Obs	4363	4358	4353	4343
adj. R-sq	0.009	0.134	0.161	0.173
Eurodollar Market	Number of Lags in Control			
	(None)	(5)	(10)	(20)
Release Dummy	0.3324*	0.2495	0.2962*	0.3019*
	(2.20)	(1.84)	(2.25)	(2.19)
No. Obs	2627	2622	2617	2607
adj. R-sq	0.001	0.125	0.137	0.146

Table 2.4: Market Reaction to the Overall Content of FOMC Minutes

This table reports the coefficient estimates for Regression (2.8). The dependent variable is the 15-minute unexpected volatility computed as the raw volatility minus the 20-day moving average, according to Equation 2.2 in the text. The independent variables are document-level tone scores computed according to Equation (2.6) without multiplying any topic proportions. The top panel reports the results using scores computed using the merged lexicon of Harvard-IV-4 and Loughran and McDonald (2011) lexicons and the bottom panel reports results using the LM lexicons only. Control variables are defined in Section 2.4.3 of the text. The estimates use 138 FOMC minutes released between 2000 and 2015.

Panel A. Unexpected Volatility; Merged Lexicon

	Net Tone	Pos Tone	Neg Tone
	(1)	(2)	(3)
Document Tone	-0.0159	-0.0086	0.0031
	(-1.69)	(-1.44)	(0.32)
Interest Rate	-0.007	-0.0085	-0.0074
	(-0.82)	(-1.32)	(-0.86)
Unemployment	-0.0062	-0.0081	-0.0067
	(-1.03)	(-1.79)	(-1.10)
Recession	-0.0486	-0.0072	-0.0286
	(-1.75)	(-0.41)	(-0.97)
N	138	138	138
R-sq	0.008	0.011	-0.012

Panel B. Unexpected Volatility; LM Lexicon

	Net Tone	Pos Tone	Neg Tone
	(1)	(2)	(3)
Document Tone	-0.012	-0.0057	0.0046
	(-1.29)	(-0.84)	(0.52)
Interest Rate	-0.006	-0.0074	-0.0074
	(-0.69)	(-1.09)	(-0.86)
Unemployment	-0.0074	-0.0083	-0.0071
	(-1.22)	(-1.83)	(-1.16)
Recession	-0.0445	-0.01	-0.0313
	(-1.55)	(-0.54)	(-1.10)
N	138	138	138
R-sq	-0.001	0.001	0.011

Table 2.5: FOMC Topic Proportion and Market Reaction

Columns (1) to (4) of this table report the coefficient estimates for Regression (2.9a). The dependent variable is the 15-minute unexpected volatility computed as the raw volatility minus the 20-day moving average, according to Equation 2.2 in the text. The independent variables are document-level proportions for each topic. Control variables are defined in Section 2.4.3 of the text. The estimates use 138 FOMC minutes released between 2000 and 2015.

	Unexpected Volatility				Directional Price Change			
	Equity Market (SPY)		Debt Market (Eurodollar)		Equity Market (SPY)		Debt Market (Eurodollar)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Policy	0.0222*	0.0256*	0.3613**	0.0256*	0.0450	0.066727*	-0.2219**	-0.1789**
	(2.16)	(2.55)	(2.62)	(2.27)	(1.56)	(2.14)	(-2.60)	(-2.58)
Inflation	0.0212*	0.0235*	0.1629	0.1938*	-0.0150	0.0059	-0.1470	-0.1361
	(2.05)	(2.30)	(1.83)	(2.00)	(-0.64)	(0.22)	(-0.72)	(-0.55)
Market	0.0045	-0.0144	0.0775	0.0146	0.0255	-0.0131	0.1878	0.1052
	(0.37)	(-0.74)	(0.11)	(0.29)	(1.13)	(-0.38)	(1.47)	(1.16)
Employment	0.0170**	0.0168**	0.2611*	0.3008**	0.0249	0.0204	-0.2915	-0.2034
	(2.70)	(2.66)	(2.35)	(2.64)	(1.29)	(1.01)	(-1.84)	(-1.61)
Trade	-0.0022	0.0017	0.0948	0.1313	-0.0178	-0.0135	0.1719	0.0125
	(-0.42)	(0.32)	(1.26)	(1.37)	(-0.89)	(-0.65)	(0.26)	(0.12)
Consumption	-0.0004	0.0048	0.1076	0.0702	0.0007	0.0114	0.0356	0.0532
	(-0.08)	(0.82)	(0.53)	(0.57)	(0.03)	(0.52)	(1.10)	(1.18)
Investment	0.0294*	0.0306*	-0.0090	0.0007	0.0383	0.0331	-0.0111	-0.0089
	(1.98)	(2.27)	(-0.25)	(0.02)	(1.29)	(1.11)	(-0.72)	(-0.43)
Control Variables								
Interest Rate		-0.0004		0.1633		0.0107		0.1218
		(-0.04)		(1.16)		(0.49)		(0.88)
Unemployment		0.0237		0.1509		0.0197		-0.1086
		(1.63)		(1.61)		(1.70)		(-0.53)
Recession		0.0152		0.3124		-0.1261		-0.0044
		(0.75)		(1.07)		(-1.22)		(-0.10)
N	138	138	88	88	138	138	88	88
adj. R-sq	0.073	0.075	0.060	0.068	0.000	0.013	0.015	0.011

Table 2.6: FOMC Topic Net Tone Score and Market Reaction

Columns (1) to (4) of this table report the coefficient estimates for Regression (2.9b). The dependent variable is the 15-minute unexpected volatility computed as the raw volatility minus the 20-day moving average, according to Equation 2.2 in the text. The estimates use 138 FOMC minutes released between 2000 and 2015.

	Unexpected Volatility				Directional Price Change			
	Equity Market (SPY)		Debt Market (Eurodollar)		Equity Market (SPY)		Debt Market (Eurodollar)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Policy	-0.0014	-0.0137	-0.3448**	-0.3512*	0.0609*	0.0670**	-0.4146**	-0.4412*
	(-0.20)	(-1.48)	(-2.58)	(-2.23)	(2.44)	(2.70)	(-2.62)	(-2.38)
Inflation	-0.0180*	-0.0187**	-0.3719***	-0.3723***	0.0443**	0.0445*	-0.4740**	-0.4839**
	(-2.34)	(-2.77)	(-3.11)	(-3.25)	(2.76)	(2.49)	(-2.58)	(-2.67)
Market	-0.0047	-0.0024	0.0656	0.1001	-0.0560**	-0.0556**	0.1018	0.0693
	(-0.69)	(-0.32)	(0.20)	(0.30)	(-2.73)	(-2.65)	(0.74)	(0.62)
Employment	-0.0155*	-0.0154*	-0.1731*	-0.2087	0.0433	0.0601*	-0.0994	-0.2330
	(-2.46)	(-2.33)	(-2.04)	(-1.86)	(1.91)	(2.36)	(-1.87)	(-1.59)
Economy	0.0279	0.0350*	0.1753	0.1386	0.0025	-0.0143	-0.1439	-0.1171
	(1.88)	(2.26)	(1.54)	(1.49)	(0.14)	(-0.82)	(-0.65)	(-1.37)
Trade	-0.0075	-0.0096	-0.4511	-0.0797	0.0030	0.0050	0.0097	0.0126
	(-0.84)	(-1.05)	(-0.12)	(-0.04)	(0.16)	(0.03)	(0.38)	(0.35)
Consumption	0.0157*	0.0144*	0.1695	0.0878	-0.0364	-0.0539*	0.3381	0.1917
	(2.00)	(2.02)	(1.43)	(1.19)	(-1.38)	(-2.03)	(0.72)	(0.52)
Investment	-0.0131	-0.0136	0.0901	-0.0478	-0.0595*	-0.0698**	-0.0015	-0.0123
	(-1.35)	(-1.34)	(0.07)	(-0.82)	(-2.55)	(-3.00)	(-0.81)	(-0.86)
Control Variables								
Interest Rate		-0.0002		0.1683		0.0195		0.1200
		(-0.98)		(0.64)		(1.07)		(0.52)
Unemployment		0.0221		0.0474		0.0169		-0.1477
		(1.69)		(1.25)		(1.05)		(-0.21)
Recession		0.0173		-0.0735		-0.1170		-0.0093
		(1.70)		(-0.14)		(-1.84)		(-0.36)
N	138	138	88	88	138	138	88	88
adj. R-sq	0.069	0.077	0.084	0.080	0.015	0.025	0.012	0.009

Table 2.7: Negative Topic Tones and Market Volatility

This table reports the coefficient estimates for Regression (2.9b), where the independent variable is the Negative Tone Scores for each of the 8 LDA-identified topics computed according to Equation (2.6) of the text. The dependent variable is the 15-minute unexpected volatility computed as the raw volatility minus the 20-day moving average, according to Equation 2.2 in the text. The equity market regression uses transaction prices of SPY to measure intraday market volatilities. The Eurodollar market regression uses 3-month LIBOR rate volatilities implied by the nearest maturity Eurodollar futures contract. The estimates use 138 FOMC minutes released between 2000 and 2015.

	Markets			
	Equity (SPY)		Debt (Eurodollar)	
	(1)	(2)	(3)	(4)
Policy	0.0189*	0.0151	0.2517**	0.2508**
	(2.28)	(1.91)	(2.43)	(2.47)
Inflation	0.0162*	0.0206**	-0.1538**	-0.1832*
	(2.53)	(2.95)	(-2.50)	(-2.18)
Market	-0.0154	-0.0187*	-0.1917*	-0.2021*
	(-1.63)	(-2.05)	(-2.15)	(-2.12)
Employment	0.0158*	0.0168*	0.0324	0.1297
	(2.10)	(2.27)	(1.32)	(0.81)
Economy	-0.0089	-0.0037	0.2994***	0.3318***
	(-1.46)	(-0.56)	(3.22)	(3.39)
Trade	-0.0016	-0.0024	0.1172	0.0453
	(-0.31)	(-0.47)	(0.80)	(0.71)
Consumption	0.0015	-0.0025	-0.1826	-0.0333
	(0.26)	(-0.46)	(-0.07)	(-0.22)
Investment	-0.0023	-0.0015	-0.0647	-0.1021
	(-0.44)	(-0.29)	(-1.35)	(-1.45)
Control Variables				
Interest Rate		-0.0007		0.2293
		(-0.17)		(1.04)
Unemployment		0.0181*		0.0837
		(2.46)		(1.16)
Recession		0.0142		0.0002
		(0.59)		(0.01)
N	138	138	88	88
adj. R-sq	0.054	0.069	0.128	0.142

Table 2.8: Positive Topic Tones and Market Volatility

This table reports the coefficient estimates for Regression (2.9b), where the independent variable is the Positive Tone Scores for each of the 8 LDA-identified topics computed according to Equation (2.6) of the text. The dependent variable is the 15-minute unexpected volatility computed as the raw volatility minus the 20-day moving average, according to Equation 2.2 in the text. The equity market regression uses transaction prices of SPY to measure intraday market volatilities. The Eurodollar market regression uses 3-month LIBOR rate volatilities implied by the nearest maturity Eurodollar futures contract. The estimates use 138 FOMC minutes released between 2000 and 2015.

	Markets			
	Equity (SPY)		Debt (Eurodollar)	
	(1)	(2)	(3)	(4)
Policy	-0.0156*	-0.0131*	-0.0012	-0.0253
	(-2.16)	(-2.13)	(-0.02)	(-0.37)
Inflation	0.0005	0.0009	-0.0678	-0.0942
	(0.07)	(0.12)	(-0.88)	(-1.13)
Market	-0.0277*	-0.0341*	-0.1553*	-0.1556*
	(-2.55)	(-2.53)	(-2.18)	(-2.11)
Employment	0.0021	0.0039	0.1740	0.1246
	(0.31)	(0.67)	(1.74)	(1.80)
Economy	-0.0293***	-0.0227**	0.1322*	0.1807*
	(-3.71)	(-2.75)	(2.06)	(2.14)
Trade	-0.0029	-0.0028	0.1007*	0.0852
	(-0.59)	(-0.59)	(2.03)	(1.79)
Consumption	0.0036	0.001	-0.0045	0.0120
	(0.53)	(0.13)	(-0.28)	(0.13)
Investment	0.007	0.0066	-0.0244	-0.0469
	(1.01)	(0.96)	(-0.52)	(-0.14)
Control Variables				
Interest Rate		-0.0005		0.1619
		(-0.21)		(1.12)
Unemployment		0.0147		0.0800
		(1.38)		(1.16)
Recession		0.0118		0.1192
		(0.51)		(0.42)
N	138	138	88	88
adj. R-sq	0.050	0.051	0.073	0.064

Table 2.9: Topic Proportion and Macroeconomic Variables

This table reports the coefficient estimates for Regression (2.11a). The independent variables are document-level proportions for each topic. The independent variables are related to macroeconomic conditions and defined in Section 2.4.3 of the text. These variables are also used as controls in other regressions. The estimates use 196 FOMC minutes released between 1990 and 2015.

	Topics							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Policy	Inflation	Market	Employment	Growth	Trade	Consumption	Investment
Interest Rate	0.0487	-0.432***	-0.221***	-0.107*	0.409***	0.00231	-0.164***	0.244***
	(1.06)	(-13.12)	(-8.96)	(-2.47)	(13.46)	(0.05)	(-3.71)	(6.38)
Unemployment	0.0690	-0.0425	0.345***	-0.156***	-0.0799*	-0.0694	-0.179***	-0.158***
	(1.42)	(-1.22)	(13.23)	(-3.38)	(-2.48)	(-1.42)	(-3.81)	(-3.90)
Recession Dummy	-0.446	-0.0926	0.974***	-0.913***	-0.102	0.227	0.153	-0.138
	(-1.85)	(-0.54)	(7.56)	(-4.00)	(-0.64)	(0.94)	(0.66)	(-0.69)
N	196	196	196	196	196	196	196	196
adj. R-sq	0.015	0.493	0.717	0.114	0.568	0.001	0.085	0.315

Table 2.10: Topic Tone Scores and Macroeconomic Variables

Panel A: Net Tone

	Topic							
	Policy	Inflation	Market	Employment	Economy	Trade	Consumption	Investment
Interest Rate	0.0036	0.1160**	-0.0129	0.1334***	-0.0029	0.0030	0.0475	0.0536
	(0.08)	(2.78)	(-0.31)	(3.66)	(-0.07)	(0.06)	(1.16)	(1.35)
Unemployment	0.0308	-0.0481	-0.1683***	-0.2332***	-0.1398**	0.0294	-0.0759	-0.0965*
	(0.65)	(-1.09)	(-3.81)	(-6.03)	(-3.15)	(0.60)	(-1.76)	(-2.30)
Recession	-0.9529***	-1.2035***	-1.2021***	-1.1137***	-1.2576***	-0.4359	-1.4928***	-1.5910***
	(-4.07)	(-5.50)	(-5.50)	(-5.83)	(-5.73)	(-1.80)	(-6.99)	(-7.67)
adj. R-sq	0.067	0.185	0.187	0.378	0.180	0.003	0.222	0.266

Panel B: Positive Tone

	Topic							
	Policy	Inflation	Market	Employment	Economy	Trade	Consumption	Investment
Interest Rate	0.0127	-0.3390***	-0.1157**	0.1125**	0.2913***	-0.0408	-0.0035	0.1746***
	(0.28)	(-8.58)	(-2.93)	(2.70)	(7.71)	(-0.89)	(-0.08)	(4.37)
Unemployment	0.1252**	-0.0944*	0.1964***	-0.1530***	-0.0916*	-0.0679	-0.1328**	-0.1758***
	(2.60)	(-2.26)	(4.70)	(-3.46)	(-2.29)	(-1.40)	(-2.81)	(-4.16)
Recession	-0.4292	-0.0616	0.9733***	-0.7880***	-0.3288	0.4367	-0.6714**	-0.5139*
	(-1.81)	(-0.30)	(4.71)	(-3.61)	(-1.66)	(1.81)	(-2.88)	(-2.46)
adj. R-sq	0.037	0.271	0.274	0.187	0.333	0.013	0.072	0.256

Panel C: Negative Tone

	Topic							
	Policy	Inflation	Market	Employment	Economy	Trade	Consumption	Investment
Interest Rate	0.0064	-0.3109***	-0.0469	-0.0553	0.3267***	-0.0269	-0.0645	0.1029*
	(0.14)	(-8.51)	(-1.28)	(-1.27)	(8.55)	(-0.59)	(-1.51)	(2.54)
Unemployment	0.0667	-0.0159	0.2367***	0.1509**	0.0349	-0.0685	-0.0252	-0.0481
	(1.37)	(-0.41)	(6.09)	(3.27)	(0.86)	(-1.42)	(-0.56)	(-1.12)
Recession	0.4631	1.0189***	1.4762***	0.6496**	0.8651***	0.6803**	1.3034***	1.5669***
	(1.93)	(5.33)	(7.68)	(2.85)	(4.32)	(2.86)	(5.84)	(7.40)
adj. R-sq	0.015	0.377	0.371	0.114	0.318	0.036	0.153	0.236

Table 2.11: Volatility Reaction to the Magnitude of Tones

This table reports the coefficient estimates for Regression (2.14a). The dependent variable is the 15-minute unexpected volatility computed as the raw volatility minus the 20-day moving average, according to Equation 2.2 in the text. The independent variables are the absolute values of document-level Net Tone scores for each of the 8 LDA-identified topics computed according to Equation (2.6). Control variables are defined in Section 2.4.3 of the text. The estimates use 138 FOMC minutes released between 2000 and 2015.

	Models	
	(1)	(2)
Policy	0.0192*	0.0155*
	(2.60)	(2.19)
Inflation	0.0133*	0.0160*
	(2.01)	(2.00)
Market	0.0153*	0.0173*
	(2.00)	(2.02)
Employment	0.0115*	0.0138*
	(2.06)	(2.19)
Economy	0.0018	0.0035
	(0.32)	(0.50)
Trade	-0.0018	0.0025
	(-0.33)	(0.51)
Consumption	0.0103	0.0049
	(1.70)	(0.68)
Investment	0.009	0.0143
	(1.25)	(1.96)
Control Variables		
Interest Rate		-0.0010
		(-0.17)
Unemployment		0.0156
		(1.77)
Recession		0.0157
		(1.34)
N	138	138
adj. R-sq	0.068	0.070

Table 2.12: Market Reaction to Tone Changes

Columns (1) and (2) of this table report the coefficient estimates for Regression (2.13a). The dependent variable is the 15-minute unexpected volatility computed as the raw volatility minus the 20-day moving average, according to Equation 2.2 in the text. Columns (3) and (4) of this table report the coefficient estimates where the dependent variable is directional change in SPY. The independent variables are the changes in document-level Net Tone scores for each of the 8 LDA-identified topics computed according to Equation (2.6). Control variables are defined in Section 2.4.3 of the text. The estimates use 138 FOMC minutes released between 2000 and 2015.

	Unexpected Volatility		Directional Price Change	
	(1)	(2)	(3)	(4)
Policy	0.0244***	0.0274***	0.0756***	0.0723***
	(5.12)	(5.21)	(5.04)	(5.22)
Inflation	-0.0094	-0.0187*	0.0579*	0.0547*
	(-1.10)	(-2.15)	(2.23)	(2.25)
Market	-0.0144*	-0.0098*	-0.0738	-0.0726*
	(-2.16)	(-2.44)	(-1.95)	(-2.05)
Employment	0.0093	0.0151	0.0403*	0.0483*
	(0.47)	(0.71)	(2.09)	(2.01)
Economy	0.0013	0.0018	0.0032	0.0026
	(1.09)	(1.30)	(0.57)	(0.43)
Trade	-0.0021	-0.0028	-0.0075	-0.0052
	(-0.24)	(-0.29)	(-0.57)	(-1.22)
Consumption	0.0024	0.0006	0.0042	-0.0006
	(0.84)	(0.17)	(0.02)	(-0.12)
Investment	-0.0009	0.0001	0.0014	0.0023
	(-0.64)	(0.09)	(0.22)	(0.24)
Control Variables				
Interest Rate		-0.0025		
		(-0.60)		(1.02)
Unemployment		0.0115		0.0119
		(1.41)		(0.93)
Recession		0.0153		-0.1401
		(1.42)		(-1.78)
N	138	138	138	138
adj. R-sq	0.058	0.077	0.014	0.027

Table 2.13: Market Reaction Pre- and Post-2011

This table reproduces Table 2.6 for different subsamples. The equity market regression uses transaction prices of SPY to measure intraday market returns and volatilities. The Eurodollar market regression uses 3-month LIBOR rate and volatilities implied by the nearest maturity Eurodollar futures contract. The pre-2011 sample uses 101 FOMC minutes released between 2000 and 2011. The post-2011 sample uses 101 FOMC minutes released between 2011 and 2015. Control variables are defined in Section 2.4.3 of the text.

	Unexpected Volatility				Directional Price Change			
	Pre-2011		Post-2011		Pre-2011		Post-2011	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Policy	-0.0123	-0.0127	-0.0228	-0.0243	0.0523*	0.0649**	0.0637*	0.0711*
	(-1.15)	(-1.31)	(-1.82)	(-1.57)	(2.24)	(2.69)	(2.50)	(2.53)
Inflation	-0.0150	-0.0128*	-0.0261*	-0.0219*	0.0644*	0.0737*	0.0388	0.0203
	(-1.55)	(-2.37)	(-2.55)	(-2.47)	(2.14)	(2.32)	(1.78)	(1.72)
Market	-0.0096	-0.0043	-0.0056	-0.0050	-0.0631*	-0.0526*	-0.0464*	-0.0741**
	(-1.62)	(-1.09)	(-0.53)	(-0.37)	(-2.49)	(-1.99)	(-2.54)	(-3.65)
Employment	-0.0194*	-0.0120*	-0.0288*	-0.0279*	0.0416	0.0632*	0.0470**	0.0642*
	(-2.11)	(-2.22)	(-2.01)	(-2.42)	(1.61)	(2.20)	(2.77)	(2.65)
Economy	-0.0158	-0.0162	0.0313*	0.0297*	0.0198	0.0202	-0.0139	-0.0088
	(-1.96)	(-1.98)	(2.50)	(2.34)	(1.14)	(0.32)	(-1.03)	(-0.72)
Trade	-0.0094	-0.0089	0.0001	-0.0018	0.0165	0.0127	-0.0206	0.0028
	(-1.41)	(-1.31)	(0.01)	(-0.18)	(0.72)	(0.54)	(-0.85)	(0.12)
Consumption	0.0317**	0.0308**	0.0144	0.0153	-0.0178	-0.0500	-0.0315	-0.0442
	(3.25)	(3.01)	(1.14)	(1.14)	(-0.21)	(-0.13)	(-0.91)	(-0.24)
Investment	0.0047	0.0024	-0.0365*	-0.0351*	-0.0645*	-0.0770*	-0.0444*	-0.0628*
	(0.50)	(0.26)	(-2.34)	(-2.34)	(-2.23)	(-2.60)	(-2.05)	(-2.14)
Interest Rate		-0.0022		0.0016		0.0071		0.1572*
		(-0.13)		(0.54)		(0.68)		(2.31)
Unemployment		0.0117		0.0114		0.0105		0.0155
		(1.62)		(0.17)		(0.58)		(0.72)
Recession		0.0179		N/A		-0.0463		N/A
		(0.71)		N/A		(-0.78)		N/A
N	101	101	37	37	101	101	37	37
adj. R-sq	0.043	0.040	0.140	0.089	0.015	0.018	0.127	0.172

Table 2.14: Statements and Minutes

Panel A of this table presents the coefficient estimates from Regression (2.9b), where the independent variable is the document-level Net Tone Scores from FOMC Statements, computed according to Equation (2.6) of the text. The Meeting Day subsample uses days when the FOMC meetings take place, and the Minutes Release Day subsample uses the days when the corresponding meeting minutes are released. Panel B of this table presents the regression estimates of the Minutes topic Net Tone Scores regressed on the document-level Net Tone Scores from FOMC statements. Control variables are defined in Section 2.4.3 of the text.

Panel A: Statement Tone and Volatility

	Raw Volatility				Unexpected Volatility			
	Meeting Day		Minutes Release Day		Meeting Day		Minutes Release Day	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Statement Net Tone	0.0612*	0.0610*	0.0117	0.0031	-0.0416*	-0.0403*	0.0123	0.0096
	(2.40)	(2.49)	(0.60)	(0.15)	(-2.03)	(-2.12)	(1.54)	(1.25)
Interest Rate		0.0257		0.0228		-0.0093		-0.0009
		(0.83)		(0.81)		(-0.96)		(-0.96)
Unemployment		0.0103		0.0187		0.0106*		0.0177*
		(1.02)		(1.82)		(2.08)		(2.05)
Recession		-0.0616		-0.0700		0.0169		0.0032
		(-0.12)		(-1.15)		(0.37)		(0.11)
adj. R-sq	0.032	0.020	-0.005	0.002	0.027	0.041	0.017	0.036

Panel B: Predictability of Minutes' Contents Using Statements

	Dependent Variable=Net Tone Score of Each Minutes Topic							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Policy	Inflation	Market	Employment	Economy	Trade	Consumption	Investment
Statement Tone	0.1021	0.2070*	-0.0640	-0.0213	0.0699	-0.0805	0.0102	-0.1342
	(1.87)	(2.20)	(-0.63)	(-0.27)	(0.89)	(-0.93)	(0.15)	(-1.44)
Interest Rate	-0.0623*	0.3682***	-0.0005	0.1325*	-0.1364	0.0071	0.0454	0.0169
	(-2.14)	(4.64)	(-0.01)	(2.44)	(-0.89)	(0.62)	(0.91)	(0.30)
Unemployment	-0.0511	0.0199	-0.1350***	-0.2093***	-0.0269*	0.0700	-0.0717	0.0097
	(-0.92)	(0.33)	(-5.97)	(-4.31)	(-2.13)	(1.04)	(-1.51)	(0.16)
Recession	-0.3741**	-0.8616***	0.9646**	-0.6085**	-1.0868***	-0.5187	-1.5546***	-1.5732***
	(-3.07)	(-4.12)	(2.92)	(-3.06)	(-4.59)	(-1.74)	(-7.83)	(-6.79)
adj. R-sq	0.093	0.269	0.183	0.394	0.150	0.007	0.269	0.240

Table 2.15: Market Reaction to Tone Computed Using Market-Weighted Lexicons

This table reproduces Table 2.6 for different Net Tone Score measures. The independent variables are document-level Net Tone scores for each of the 8 LDA-identified topics, computed according to Equation (2.6), using the market-weighted lexicon developed by Jegadeesh and Wu (2013). Control variables are defined in Section 2.4.3 of the text. The estimates use 138 FOMC minutes released between 2000 and 2015.

	Unexpected Volatility		Directional Price Change	
	(1)	(2)	(3)	(4)
Policy	-0.003	-0.0073*	0.0531	0.05439*
	(-0.99)	(-2.11)	(1.92)	(2.10)
Inflation	-0.0097*	-0.0135*	0.0061	0.0133
	(-2.05)	(-2.54)	(0.29)	(0.60)
Market	-0.0045	-0.0127	-0.0399*	-0.0474**
	(-0.98)	(-1.96)	(-2.13)	(-2.94)
Employment	-0.0027	-0.0019	0.0303	0.0597*
	(-0.52)	(-0.36)	(1.72)	(2.12)
Economy	0.0149**	0.0188**	0.0828**	0.0754**
	(2.77)	(3.13)	(3.16)	(3.10)
Trade	-0.0183	-0.0142	0.0461	0.0519
	(-1.57)	(-1.26)	(1.35)	(1.45)
Consumption	-0.0191*	-0.0240**	-0.0199	-0.0086
	(-2.58)	(-2.94)	(-0.79)	(-0.33)
Investment	0.0316***	0.0343***	-0.0707*	-0.0709*
	(3.48)	(3.92)	(-2.25)	(-2.25)
Control Variables				
Interest Rate		-0.0013		0.0191
		(-0.70)		(1.40)
Unemployment		0.0134**		0.0252
		(2.11)		(0.65)
Recession		0.0153		-0.1241
		(0.60)		(-0.90)
N	138	138	138	138
adj. R-sq	0.067	0.077	0.027	0.033

APPENDIX

A.1. Appendix for Chapter 1

A.1.1. List of Variables and Their Construction

- $Size$: Market capitalization of firms as of December of the previous calendar year. Computed as (PRC×SHROUT) with both variables obtained from CRSP (domestic firms) and Compustat Daily Price data (international firms).

- BM: Ratio of book equity to market capitalization as of December of the previous calendar year.

- PE: Price adjusted for splits (PRC/CFACPR) divided by adjusted earnings per share (EPSPX/AJEX).

- ROA: Net profits after taxes (NI) plus interest expenses (XINT) divided by total assets (AT) as of December of the previous calendar year.

- Lev: Long-term liabilities (DLTT) plus short-term liabilities (DLC) divided by market capitalization as of December of the previous calendar year.

- $INVL$: Total inventory (INVTQ) divided by total assets (ATQ) in quarter $t-1$.

- GM: Gross margin, computed as revenue (SALEQ) minus cost of goods sold (COGSQ) divided by revenue in quarter t.

- $CAPEX$: Capital expenditure (CAPXQ) divided by total assets (ATQ) in quarter $t-1$.

- RD: Research and development expenditure (XRDQ) divided by total assets (ATQ) in quarter $t-1$.

- $Cash$: Cash and short-term investments (CHEQ) divided by total assets (ATQ) in quarter $t-1$.

- RE: Retained earnings (REQ) divided by total assets (ATQ) in quarter $t-1$.

- NEI: Net equity issuance, computed as sale of common and preferred stock (SSTKQ) minus purchase of common and preferred stock (PRSTKCQ), divided by total assets (ATQ) in quarter $t-1$.

- NDI: Net debt issuance, computed as long-term debt issuance (DLTISQ) minus long-term debt reduction (DLTRQ), divided by total assets (ATQ) in quarter $t-1$.

- D^n: Dummy variable that equals to 1 if one of the firm's supplier at a distance of n connections experiences a shock.

- γ: Supply relationship share, computed as $\gamma_{ji,t} = \dfrac{V_{ji,t}}{COGS_{i,t}}$.

- MP: Market power measure computed as $MP_{i,t} = \dfrac{\text{Size}_{i,t-1}}{\sum_k^{N_i} \text{Size}_{k,t-1}}$, where $Size$ is defined above.

A.1.2. Data Construction Methodologies

The LDA Algorithm

Prior to the advent of probabilistic topic models, the classification of textual documents and the inference of their contexts are done either manually or in a static fashion using word-based approaches such as keyword searches or latent semantic analysis. Probabilistic topic models such as the Latent Dirichlet Allocation (LDA) algorithm remove this limitation and allow for automated and accurate classification of documents on a large, "big data" scale. First developed by Blei et al. (2003), the LDA belongs to a broader class of probabilistic topic models that use hierarchical Bayesian analysis to uncover the underlying semantic structure of textual documents. The advantage of this approach is discussed in the main text. Here I first illustrate the approach with a simple example. Suppose that the full vocabulary used in firm disclosures consists of only $V = 4$ words (ignore common words such as *I*, *the*, etc): {*earthquake*, *demolish*, *economy*, *consumption*}. Suppose there are $D = 3$ disclosures:

1. *An earthquake demolished our factory.*

2. *We are unable to meet consumer demand due to strong economy.*

3. *An earthquake demolished our factory. In addition consumer demand is very strong due to the economy, so we are unable to meet the demands.*

A human reader would intuitively recognize that the first document is primarily in the context of natural disasters and the second is about an economy-driven demand shock. The third document is a mixture of both. Suppose I fit the LDA model with $N = 2$ topics. If the model performs satisfactorily, then first, the posterior topic distributions should clearly and intuitively identify the topics and thus be something similar to:

- $\hat{\beta}_1 \equiv \{\hat{P}_{topic1}(earthquake), \hat{P}_{topic1}(demolish), \hat{P}_{topic1}(economy), \hat{P}_{topic1}(consumer)\}$

  $= \{0.55, 0.43, 0.01, 0.01\}$

- $\hat{\beta}_2 \equiv \{\hat{P}_{topic2}(earthquake), \hat{P}_{topic1}(demolish), \hat{P}_{topic1}(economy), \hat{P}_{topic1}(consumer)\}$

  $= \{0.01, 0.01, 0.60, 0.48\}$

Next, the posterior topic mixture in each document should correspond to the human reader's intuition:

- $\hat{\theta}_1 \equiv \{\hat{P}_{document1}(Topic1), \hat{P}_{document1}(Topic2)\} = \{0.99, 0.01\}$

- $\hat{\theta}_2 \equiv \{\hat{P}_{document2}(Topic1), \hat{P}_{document2}(Topic2)\} = \{0.01, 0.99\}$

- $\hat{\theta}_3 \equiv \{\hat{P}_{document3}(Topic1), \hat{P}_{document3}(Topic2)\} = \{0.51, 0.49\}$

I proceed with my LDA classification of firm disclosures by generalizing this example to the sample of $D = 19,771$ disclosures. Stop words, location and industry-specific terms, and other commonly appearing words, such as *a, the*, etc., are removed prior to processing. This results in a collection of $V = 53,971$ English words.[1]

I hypothesize that there are $N = 20$ unique topics in the document. Here, each of the N topics represents a distribution over the V words in the disclosure vocabulary, and each document is a mixture of the N topics. I assume that the observable data, i.e. words in each document, is generated from a probabilistic data generating process parameterized as follows:

1. Each of document $d = 1, \ldots, D$ contains a mixture of N topics. Let the proportion of topic n in document d be $\theta_{d,n}$ and let the vector $\theta_d = [\theta_{d,1}, \ldots, \theta_{d,N}]'$ represent the true topic mixture of document d. For each d, I assume that this mixture follows an order-N Dirichlet distribution over the N topics, governed by the latent, parameter vector μ of size N:

$$\theta_d \sim Dirichlet_N(\mu)$$

2. Given document d's topic mixture θ_d, let the assignment of each word i in document d into topics be $Z_{d,i}$, where $Z_{d,i} \in \{1, \ldots, N\}$. I assume that this assignment follows the multinomial distribution governed by the document-specific topic vector θ_d described in the previous step:

$$Z_{d,i}|\theta_d \sim Multinomial(\theta_d) \tag{A.1}$$

   Suppose there are I_d unique words in document d. Let the vector Z_d denote the collection of the topic assignment of all words within d, i.e. $Z_d = \{Z_{d,i}\}_{i=1}^{I_d}$

3. The N topic distributions (applied universally to all documents) are in the collection $\beta = \{\beta_1, \ldots, \beta_N\}$. Each topic β_n also follows an order-V Dirichlet distribution over

[1] I do not stem the words before processing i.e. each inflection of a word is treated as a new word.

the V words, governed by the latent scalar parameter ϕ:

$$\beta_n \sim Dirichlet_V(\phi) \tag{A.2}$$

4. For each word i in document d, there are V choices to choose from the disclosure vocabulary. Conditional on the chosen topic for word i in Step 2 above (i.e. a draw from Distribution (A.1)), and on the structure of the topic distribution from Step 3 (i.e. a draw from Distribution (A.2)), I assume that actual choice of the word, $W_{d,i}$, follows a multinomial distribution governed by the resulting word-topic assignment $\beta_{Z_{d,i}}$:

$$W_{d,i}| \left(\{\beta_n\}_{n=1}^N, Z_{d,i}\right) \sim Multinomial(\beta_{Z_{d,i}})$$

   Similarly, let the W_d denote the collection of the vocabulary choice of all words within document d: $W_d = \{W_{d,i}\}_{i=1}^{I_d}$

The above four distributions constitute the latent data generating process that results in my observable document collection $\{W_d\}_{d=1}^D$. Recall that they are not directly observable to the researcher. Instead, the only observable data is the occurrence of the actual words i in each document d, i.e. W_d. I can then write the overall data generating process as the joint distribution of latent variables $\{\beta_n\}_{n=1}^N, \{\theta_d\}_{d=1}^D, \{Z_d\}_{d=1}^D$ and the observable variable $\{W_d\}_{d=1}^D$:

$$\begin{aligned}
&P\left(\{\beta_n\}_{n=1}^N, \{\theta_d\}_{d=1}^D, \{Z_d\}_{d=1}^D, \{W_d\}_{d=1}^D\right) \\
&= \prod_{n=1}^N P(\beta_n) \prod_{d=1}^D P(\theta_d) \left[\prod_{i=1}^{I_d} P(Z_{d,i}|\theta_d) P\left(W_{d,i}|\{\beta_n\}_{n=1}^N, Z_{d,i}\right)\right]
\end{aligned}$$

where $P(\cdot)$ are the respective (Dirichlet or multinomial) density functions specified above.

Now that I observe my firm disclosure collection $\{W_d\}_{d=1}^D$, I can compute the posterior distribution of the document-topic structure given the observed documents using Bayes'

Rule:

$$P\left(\{\beta_n\}_{n=1}^N, \{\theta_d\}_{d=1}^D, \{Z_d\}_{d=1}^D | \{W_d\}_{d=1}^D\right) = \frac{P\left(\{\beta_n\}_{n=1}^N, \{\theta_d\}_{d=1}^D, \{Z_d\}_{d=1}^D, \{W_d\}_{d=1}^D\right)}{P\left(\{W_d\}_{d=1}^D\right)}. \tag{A.3}$$

Similar to other Bayesian inference methods, the numerator in Equation (A.3) can be easily computed. The denominator is by construction a double integral and therefore cannot be feasibly computed. However, it can be efficiently approximated using a Gibbs sampler. I use a customized Gibbs sampler written in C++ for fast implementation.

Once the posterior probabilities are computed, I compute the posterior expectations of two key latent variables, which represent the main output from the LDA algorithm:

1. Posterior vocabulary distribution for each topic: $\{\hat{\beta}_1, \ldots, \hat{\beta}_N\}$

2. Posterior topic mixture for each document in my collection: $\{\hat{\theta}_1, \ldots, \hat{\theta}_D\}$

The first set of output from my LDA procedure identifies the topics. For each topic k, $\hat{\beta}_k = [\hat{\beta}_{k,1}, \ldots, \hat{\beta}_{k,V}]'$, and each entry $\hat{\beta}_{k,j}$ represents the *probability that the word j characterizes topic k*. My whole collection of 8-K and other disclosures has $V = 72,442$ unique terms. As a result, each $\hat{\beta}_k$ contains 72,442 entries, the majority of which receives a weight close to zero. The top 5 keywords for each topic are reported in Table 1.2. The main text interprets of these topics.

The second set of output is the collection of document-level topic mixture vectors, $\{\hat{\theta}_1, \ldots, \hat{\theta}_D\}$. From this collection, each document d has one mixture, $\hat{\theta}_d = [\hat{\theta}_{d,1}, \ldots, \hat{\theta}_{d,N}]'$. Because there are 20 topics, each vector $\hat{\theta}_d$ has 20 entries, where each $\hat{\theta}_{d,n}$ corresponds to the *proportion* of document d that is devoted to topic n. The 20 entries sum up to one for each document. I keep the disclosure only if it contains more than 95% of a single topic. The main text then discusses the contexts of these topics and that these topics help identify shocks that are idiosyncratic in causes.

Network Data

I start at the year 1994, when computerized filing records become available in SEC's EDGAR system. For each firm pair ji, I identify them as in a supply relationship if either firm discloses the relationship in a filing, or either firm is identified in the other sources. The relationship is removed if either party discloses its termination. Note that in this setting, I do not capture the value of the relationships. Therefore, if supplier j and customer i are in a relationship at t, I set the relationship dummy $\gamma_{ji,t} = 1$ and 0 otherwise. A smaller portion of my sample (e.g. those identified by Bloomberg or the import-export data) does report relations with specific dollar values. In this case where firm j supplies $V_{ji,t}$ worth of goods to firm i in year t, I set $\gamma_{ji,t} = \frac{V_{ji,t}}{COGS_{i,t-1}}$. This restricted sample is useful in several robustness checks.

In essence, the existence (and for a smaller sample, the magnitude) of supply chain linkages between firms are captured by the relationship indicator γ_{ji}. For my sample of $N = 10,930$ firms, I organize all γ_{ij} parameters in the following N-by-N matrix:

$$\Gamma_t = \begin{bmatrix} \gamma_{11,t} & \cdots & \gamma_{1N,t} \\ \vdots & \ddots & \vdots \\ \gamma_{N1,t} & \cdots & \gamma_{NN,t} \end{bmatrix} \tag{A.4}$$

In this matrix, if firms i and j have no direct customer-supplier relationship, $\gamma_{ij} = \gamma_{ji} = 0$. Otherwise, the non-zero entries in Γ represent direct customer-supplier relationships, or supply chain linkages. Note that Γ needs not be symmetric. For example, if firm j is a supplier of i, but does not purchase any input from i, then γ_{ji} is nonzero while γ_{ij} is zero. In the smaller sample where $\gamma_{ji} < 1$, the sum of the ith row conveys the significance of firm i to the economy, while the sum of ith column measures firm i's degree of reliance on external intermediate goods. In this case each column sums to less than one by construction.[2]

[2]This ensures that Γ' has positive eigenvalues and the inverse, $(I - \Gamma')^{-1}$, is not singular. I use this inverse extensively in a related paper to compute the centrality of each firm.

The relationships in this economy, summarized in Γ, constitute a directed *network* where each firm i is a *node* and each relationship γ_{ij}, if nonzero, is a *link* that points from firm i (supplier) to firm j (customer). Equivalently, the supply chain network can be visualized as a directed graph where Γ represents the *adjacency matrix* of this graph. The main text then discusses the summary statistics of this network.

Finally, an important reason to represent the linkages in a network structure is to gauge the effect of supply chain linkages *beyond the immediate connections.* To illustrate this, Figure 1.1 presents a visualization of a portion of the network: 400 select US firms from technology-related industries (Fama-French industry codes 35 to 37) in the years of 2002 and 2015, respectively. There are few isolated nodes within the network. For many firms, the supply chain can be quite long, even extending to fourth- and fifth-tier suppliers. These firms will be subject to additional shocks if shocks can spill over beyond the closest connections into customers further downstream. I empirically document these spillovers in Section 1.3.

A.1.3. Additional Event Text Examples

This section lists several examples for each category of shocks characterized by the LDA algorithm. Category 1 consists of events caused by economy- or industry-wide systematic factors. Category 2 consists of events caused by potentially idiosyncratic factors but could also caused by industry-wide issues. Category 3 consists of events caused by idiosyncratic factors.

- **Example 1A, Economic Issues:** *[The firm] announced its plan to cancel the development of...power plant that it had planned to construct...because of...reduced customer demand for electricity due to the recession and slow economic recovery, surplus generating capacity in the Midwest market, and lower natural gas prices linked to expanded shale gas supplies.* (CMS Energy Corp 8-K)

- **Example 1B, Industry Issues:** *Ford Motor Co. is scrambling to find enough steel frames to keep up with demand...frame's main supplier...was having trouble building enough of the parts to keep pace with production needs....Ford has had to cancel planned overtime at the plants and has temporarily halted assembly lines during regular shifts as workers waited for more frames to arrive.* (Wall Street Journal)

- **Example 2A, Labor Issues:** *Yue Yuen Industrial Holdings Ltd., which makes shoes for Adidas AG and Nike Inc., said production...was disrupted as workers upset with organizational changes went* *on strike. About 2,000 employees...were affected...The strike cost the company \$27 million in direct costs, including lost profits and additional air freight costs.* (Reuters)

- **Example 2B, QC Issues:** *As part of ongoing quality assurance, Intel Corporation has discovered a design issue in a recently released support chip..Intel has stopped shipment of the affected support chip from its factories..and expects full volume recovery in April..Intel expects this issue to reduce revenue by approximately \$300 million as the company discontinues production.* (Intel 8-K)

- **Example 3B, Natural Disasters Hitting Supplier:** *As a result of the earthquake...supplier suspended manufacturing operations at the factory where these materials are produced...We currently have inventory of these materials...through May 24, 2011...However, many of the factors in this situation are beyond our control, and an unfavorable development relating to any of these factors could have a material adverse effect on our results of operations.* (A123 Systems 8-K)

- **Example 3C, Manmade Disasters:** *[A] blaze occurred Sept. 4 during the installation of equipment at a factory...for SK Hynix...It will take at least half a year before SK Hynix's damaged clean room is fully rebuilt...if...production is halted for more than a week, substantial shortages could lead to higher prices, benefiting all memory-chip manufacturers...customers include Apple, Samsung, Lenovo Group Ltd., Dell Inc. and Sony Corp.* (Bloomberg)

- **Example 3D, Production Disruptions:** *Rio Tinto Alcan's Laterriere Works aluminium smelter in Quebec suffered a significant power outage yesterday...leaving the plant without the adequate energy required to continue operating at full capacity...one of two production lines has been suspended...in the coming weeks, Rio Tinto Alcan will mobilise the necessary resources to restore the suspended line.* (Rio Tinto Press Release)

- **Example 3E, Adoption Failures:** *[The] Company has experienced a backlog of orders with a significant contract manufacturer in China...while the Company had thought such manufacturing delays would be temporary, the Company learned recently that this supplier had ceased manufacturing products for the Company due to difficulties the supplier is experiencing...The timing for resuming production of the products previously manufactured by this supplier is uncertain, and will depend to a significant extent on whether the Company is able to obtain certain custom tooling used by the supplier to manufacture the Company's products.* (Loud Technologies 8-K)

A.1.4. Additional Extensions and Robustness Tests

See the Online Appendix for the complete list of additional robustness checks. The first issue is to check whether firms endogenous decisions to enter into their network positions (and inventory) introduce any bias on the interaction regressions such as Regression (1.5a) and (1.5b). To do so, I exploit exogeneous variations in inventory determined by the length of *lead time*, which is in turn determined crucially by 1) distance to the supplier and 2) mode of transportation (land vs. sea). I also use two alternative measures of supplier substitutability derived from Giannetti et al. (2011) and also used by Barrot and Sauvagnat (2014). These measures are based on industry rather than firm characteristics. I find similar results with these alternative measures.

Next, because the distribution of network linkages is uneven, linear regressions could result in the observation of "spillover effects" purely due to taking averages on uneven links. This is ruled out by computing the average degree of asymmetry at each distance from the shock, then showing that links with high values are not necessarily followed by more links with high values. The linearity issue is further ruled out with similar results from a matching approach. In another test, I use only natural disasters reported by US National Oceanic & Atmospheric Administration (NOAA) in the suppliers' geographical locations as an alternative source of shocks, and find broadly similar results.

A.2. Appendix for Chapter 2

A.2.1. Part 1. Single Topic Examples

Example 1. (99% growth mandate, other topics negligible)

With regard to developments and prospects in key sectors of the economy, members noted that despite further survey indications of eroding consumer confidence, consumer expenditures had strengthened in recent months after a pause earlier in the year. The pickup had featured rising sales of motor vehicles, and while the latter had slipped recently, a number of special factors such as shortages of popular models at the end of the model year and the effects of flooding in some parts of the Midwest suggested the need to withhold judgment on any downward shift in the underlying demand for motor vehicles. Tourism was reported to have strengthened considerably in many areas this summer, though there were major exceptions. As had been true for an extended period, consumer attitudes continued to be inhibited by concerns about employment opportunities, especially given further reductions in defense spending, the ongoing restructuring and related downsizing of many business operations, and the continuing efforts by business firms to limit the number of their permanent employees in order to hold down the rising costs of health care and other nonwage worker benefits. Members noted, however, that the growth in employment thus far this year, while tending to involve many low paying jobs, had greatly exceeded the rate of expansion in 1992. In the view of at least some members, appreciable further growth was likely as business firms found it increasingly difficult in an expanding economy to meet growing demands through outsourcing, temporary workers, and overtime work. Some members also noted that the newly legislated taxes on higher incomes would tend to curtail some consumer spending. The timing of that effect was uncertain; tax liabilities had already risen, but some payments on the added tax liabilities were not due until April of 1994 and 1995.

Example 2. Inflation Mandate (99% inflation mandate, other topics negligible)

The core consumer price index advanced at a faster rate in the first quarter than it had in the fourth quarter, reflecting the pass-through of higher energy prices and a leveling off of goods prices after sizable declines last year. The higher goods price inflation owed, in part, to the recent run-up in the prices of non-oil imports, energy, and other commodities. The price index for core personal consumption expenditures also rose at a faster rate in the first quarter than it had late last year. Despite the rise in inflation this year, however, the cumulative increase in the overall consumer price index for the year ending in March was somewhat less than the advance for the twelve months ending in March 2003. In the year ending in March, the increase in the price index for total personal consumption expenditures was similar to that of a year earlier. Survey measures of near-term inflation expectations edged up somewhat in March and April, but measures of longer-term expectations decreased. With regard to labor costs, average hourly earnings of production or nonsupervisory workers on private nonfarm payrolls rose notably less for the twelve months ending in March than they had in the year-earlier period. The overall increase in the employment cost index for private industry for the twelve months ending in March was about the same as that for the twelve-month period ending a year earlier, as wages and salaries decelerated and benefits accelerated.

Example 3. Financial Market Mandate (99% market mandate, other topics negligible)

Participants noted that financial markets were volatile over the intermeeting period, as investors responded to news on the European fiscal situation and the negotiations regarding the debt ceiling in the United States. However, the broad declines in stock prices and interest rates over the intermeeting period were seen as mostly reflecting the incoming data pointing to a weaker outlook for growth both in the United States and globally as well as a reduced willingness of investors to bear risk in light of the greater uncertainty about the outlook. While conditions in funding markets had tightened, it was noted that the condition of U.S. banks had strengthened in recent quarters and that the credit quality of both businesses and households had continued to improve.

Example 4. Policy Mandate (99% policy mandate, other topics negligible)

Participants discussed a number of policy tools that the Committee might employ if it decided to provide additional monetary accommodation to support a stronger economic recovery in a context of price stability. One of the policy options discussed was an extension of the period over which the Committee expected to maintain its target range for the federal funds rate at 0 to 1/4 percent. It was noted that such an extension might be particularly effective if done in conjunction with a statement indicating that a highly accommodative stance of monetary policy was likely to be maintained even as the recovery progressed. Given the uncertainty attending the economic outlook, a few participants questioned whether the conditionality of the forward guidance was sufficiently clear, and they suggested that the Committee should consider replacing the calendar date with guidance that was linked more directly to the economic factors that the Committee would consider in deciding to raise its target for the federal funds rate, or omit the forward guidance language entirely.

A.2.2. Part 2. Multiple Topic Examples

Example 5. (56% growth, 43% inflation)

The information reviewed at this meeting suggested that economic activity had weakened further in the opening months of the year. Production cutbacks were evident in a wide range of industries, and private payrolls had fallen markedly, especially in the goods producing sector. On the positive side, consumer confidence had rebounded sharply since the cease-fire in the Persian Gulf, retail sales and housing starts had strengthened recently, and exports had continued to expand. Broad measures of prices had slowed or contracted in January and February, but excluding energy and food prices, increases in those measures were higher than in previous months. Wage increases had moderated over the past several months.

Example 6. (83% financial market, 17% policy)

Committee members and Board members agreed that, with few exceptions, the functioning of most financial markets, including interbank markets, no longer showed significant impairment. Accordingly they agreed that the statement to be released following the meeting would indicate that the Federal Reserve would be closing the Asset-Backed Commercial Paper Money Market Mutual Fund Liquidity Facility, the Commercial Paper Funding Facility, the Primary Dealer Credit Facility, and the Term Securities Lending Facility on February 1, 2010. Committee members also agreed to announce that temporary liquidity swap arrangements between the Federal Reserve and other central banks would expire on February 1. In addition, the statement would say that amounts available through the Term Auction Facility would be scaled back further, with $50 billion of 28-day credit to be offered on February 8 and 25$ billion of 28-day credit to be offered at the final auction of March 8. The statement also would note that the anticipated expiration dates for the Term Asset-Backed Securities Loan Facility remained June 30, 2010, for loans backed by new-issue commercial mortgage-backed securities, and March 31, 2010, for loans backed by all other types of collateral. Members emphasized that they were prepared to modify these plans if necessary to support financial stability and economic growth.

Example 7. (34% growth, 31% financial market, 35% policy)

Open market operations during the intermeeting period continued to be directed toward maintaining the existing degree of pressure on reserve positions. The federal funds rate rose briefly in response to year-end pressures, but it otherwise tended to remain close to the 5-1/4 percent level expected with an unchanged policy stance. Other short-term interest rates generally were unchanged to slightly higher over the intermeeting period. Rates on intermediate- and long-term securities edged higher on balance in reaction to incoming data on economic activity that were on the firm side of market expectations; the increases in such rates appeared to be tempered, however, by favorable market reactions to new data on wages and prices. The generally positive news on economic growth and inflation along with favorable reports on earnings appeared to reinforce the optimism of equity market investors, and major indexes of stock prices increased markedly further over the intermeeting period.

Example 8. (39% growth, 13% inflation, 20% financial market, 26% policy)

In their discussion of the economic situation and outlook, FOMC meeting participants indicated that the worsening financial situation, the slowdown in growth abroad, and incoming information on economic activity had led them to mark down significantly their outlook for growth. While economic activity had evidently already been slowing over the summer, the turmoil in recent weeks had apparently resulted in tighter financial conditions and greater uncertainty among businesses and households about economic prospects, further limiting their ability and willingness to make significant spending commitments. Recent measures of business and consumer sentiment had fallen to historical lows. Participants generally expected the economy to contract moderately in the second half of 2008 and the first half of 2009, and agreed that the downside risks to growth had increased. While some expected an improving financial situation to contribute to a recovery in growth by mid-2009, others judged that the period of economic weakness could persist for some time. Several participants indicated that they expected some fiscal stimulus in coming quarters, but they were uncertain about the extent and duration of the resulting support to economic activity. Participants agreed that in coming quarters inflation was likely to move down to levels consistent with price stability, reflecting the recent declines in the prices of energy and other commodities, the appreciation of the dollar, and the expected widening of margins of resource slack. Indeed, some saw a risk that over time inflation could fall below levels consistent with the Federal Reserve's dual objectives of price stability and maximum employment.

BIBLIOGRAPHY

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

D. Acemoglu, A. Ozdaglar, and A. Tahbaz-Salehi. *Networks, Shocks, and Systemic Risk*. Working Paper, Massechusettes Institute of Technology, 2015.

G. Afonso, A. Kovner, and A. Schoar. Trading partners in the interbank lending market. *Working Paper, Massachusetts Institute of Technology*, 2015.

K. Ahern. *Network centrality and the cross section of stock returns*. Working Paper, University of Southern California, 2012.

A. Ang, J. Boivin, S. Dong, and R. Loo-Kung. Monetary policy shifts and the term structure. 78(2):429–457, 2002.

E. Ang, D. A. Iancu, R. Swinney, et al. Disruption risk and optimal sourcing in multi-tier supply networks. *Working Paper, Stanford University*, 2014.

R. Anupindi and R. Akella. Diversification under supply uncertainty. *Management Science*, 39(8):944–963, 1993.

E. Bajo, T. Chemmaunr, K. Simonyan, and H. Tehranian. Underwriter networks, investor attention, and initial public offerings. *Journal of Financial Economics, Forthcoming*, 2015.

S. R. Baker, N. Bloom, and S. J. Davis. *Measuring economic policy uncertainty*. NBER Working Paper No. 21633, 2015.

S. Banerjee, S. Dasgupta, and Y. Kim. Buyer–supplier relationships and the stakeholder theory of capital structure. *The Journal of Finance*, 63(5):2507–2552, 2008.

J.-N. Barrot and J. Sauvagnat. *Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks*. Working Paper, Massechusettes Institute of Technology, 2014.

N. Bergman, R. Iyer, and R. Thakor. Financial accelerator at work: Evidence from corn fields. *Working Paper, Massachusetts Institute of Technology*, 2015.

M. Billio, M. Getmansky, A. W. Lo, and L. Pelizzon. Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics*, 104(3):535–559, 2012.

D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *Journal of machine Learning research*, 3:993–1022, 2003.

C. Boehm, A. Flaaen, and N. Pandalai-Nayar. *Input Linkages and the Transmission of*

Shocks: Firm-Level Evidence from the 2011 Tohoku Earthquake. Working Paper, Fed Board of Governors, 2014.

E. Boukus and J. V. Rosenberg. *The information content of FOMC minutes.* Working paper, Federal Reserve Bank of New York, New York, NY, 2006.

J. Y. Campbell, C. Pflueger, and L. M. Viceira. *Monetary policy drivers of bond and equity risks.* Working paper, Harvard University, Cambridge, MA, 2015.

L. Cohen and A. Frazzini. Economic links and predictable returns. *Journal of Finance*, 63 (4):1977–2011, 2008.

L. Cohen and D. Lou. Complicated firms. *Journal of financial economics*, 104(2):383–400, 2012.

L. Cohen, A. Frazzini, and C. Malloy. The small world of investing: Board connections and mutual fund returns. *Journal of Political Economy*, 116:951–979, 2008.

L. Cohen, A. Frazzini, and C. Malloy. Sell-side school ties. *The Journal of Finance*, 65(4): 1409–1437, 2010.

E. Cohen-Cole, A. Kirilenko, and E. Patacchini. Trading networks and liquidity provision. *Journal of Financial Economics*, 113(2):235–251, 2014.

O. Dessaint and A. Matray. *Do Managers Overreact to Salient Risks? Evidence from Hurricane Strikes.* Working Paper, University of Toronto, 2015.

C. E. Fee, C. J. Hadlock, and S. Thomas. Corporate equity ownership and the governance of product market relationships. *The Journal of Finance*, 61(3):1217–1251, 2006.

M. J. Fleming and M. Piazzesi. *Monetary policy tick-by-tick.* Working paper, University of Chicago, Chicago, IL, 2005.

A. T. Foerster, P.-D. G. Sarte, and M. W. Watson. Sectoral vs. aggregate shocks: A structural factor analysis of industrial production. *Journal of Political Economy*, 119(1): 1–38, 2011.

T. Foucault and L. Fresard. Learning from peers' stock prices and corporate investment. *Journal of Financial Economics*, 111(3):554–577, 2014.

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

M. Giannetti, M. Burkart, and T. Ellingsen. What you sell is what you lend? explaining trade credit contracts. *Review of Financial Studies*, 24(4):1261–1298, 2011.

J. d. Giovanni, A. Levchenko, and I. Méjean. Firms, destinations, and aggregate fluctuations. *Econometrica*, 82(4):1303–1340, 2014.

F. Gourio. Estimating firm-level risk. *Working Paper, Federal Reserve Bank of Chicago*, 2008.

K. W. Hanley and G. Hoberg. The information content of IPO prospectuses. *Review of Financial Studies*, 23(7):2821–2864, 2010.

K. B. Hendricks and V. R. Singhal. Association between supply chain glitches and operating performance. *Management science*, 51(5):695–711, 2005.

C. A. Hennessy and D. Livdan. Debt, bargaining, and credibility in firm–supplier relationships. *Journal of Financial Economics*, 93(3):382–399, 2009.

M. G. Hertzel, Z. Li, M. S. Officer, and K. J. Rodgers. Inter-firm linkages and the wealth effects of financial distress along the supply chain. *Journal of Financial Economics*, 87 (2):374–387, 2008.

G. Hoberg and G. Phillips. Text-based network industries and endogenous product differentiation. *Journal of Political Economy, Forthcoming*, 2015.

A. Hortaçsu, G. Matvos, C. Syverson, and S. Venkataraman. Indirect costs of financial distress in durable goods industries: The case of auto manufacturers. *Review of Financial Studies*, pages 1248–1290, 2013.

N. Jegadeesh and D. Wu. Word power: A new approach for content analysis. *Journal of Financial Economics*, 110(3):712–729, 2013.

N. Jegadeesh and D. Wu. *Deciphering Fedspeak: The Information Content of FOMC Minutes.* Working Paper, University of Pennsylvania, 2015.

B. Kelly, H. Lustig, and S. Van Nieuwerburgh. *Firm Volatility in Granular Networks.* Working Paper, University of Chicago, 2013.

M. Kolay and M. L. Lemmon. Spillover effects in the supply chain: Evidence from chapter 11 filings. *Working Paper, University of Utah*, 2011.

M. T. Leary and M. R. Roberts. Do peer firms affect corporate financial policy? *The Journal of Finance*, 69(1):139–178, 2014.

T. Loughran and B. McDonald. When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *Journal of Finance*, 66:35–65, 2011.

D. O. Lucca and E. Moench. The pre-fomc announcement drift. *Journal of Finance*, 70(1): 329–371, 2015.

G. Maturana and J. Nickerson. *Teachers Teaching Teachers: The Role of Networks on Financial Decisions.* Working Paper, Emory University, 2015.

L. Menzly and O. Ozbas. Market segmentation and cross-predictability of returns. *Journal of Finance*, 65(4):1555–1580, 2010.

V. Poll, N. Stoffman, and S. Yonker. The people in your neighborhood: Social interactions and mutual fund portfolios. *The Journal of Finance*, 70(6):2679–2732, 2015.

J. D. Rauh and A. Sufi. Explaining corporate capital structure: Product markets, leases, and asset similarity. *Review of Finance*, 16(1):115–155, 2012.

C. D. Romer and D. H. Romer. Does monetary policy matter? a new test in the spirit of friedman and schwartz. In *NBER Macroeconomics Annual 1989, Volume 4*, pages 121–184. MIT Press, 1989.

W. Schmidt and A. Raman. *When Supply-Chain Disruptions Matter*. Working Paper 13-006, Harvard Business School, 2013.

A. Schoar. Effects of corporate diversification on productivity. *The Journal of Finance*, 57 (6):2379–2403, 2002.

K. Shue. Executive networks and firm policies: Evidence from the random assignment of mba peers. *Review of Financial Studies*, 26(6):1401–1442, 2013.

C. A. Sims and T. Zha. Were there regime switches in U.S. monetary policy? *American Economic Review*, 96(1):54–81, 2006.

I. A. Strebulaev and T. M. Whited. Dynamic models and structural estimation in corporate finance. *Foundations and Trends in Finance*, 6(1-2), 2011.

P. C. Tetlock. Giving content to investor sentiment: The role of media in the stock market. *Journal of Finance*, 62:1139–1168, 2007.

S. Titman and R. Wessels. The determinants of capital structure choice. *The Journal of finance*, 43(1):1–19, 1988.

B. Tomlin. On the value of mitigation and contingency strategies for managing supply chain disruption risks. *Management Science*, 52(5):639–657, 2006.

D. Zhou. *The Blame Game*. Working Paper, Emory University, 2014.