Essays in Financial Economics

Lukas Felix Fischer

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

COLUMBIA UNIVERSITY

2024

Abstract

Essays in Financial Economics

Lukas Felix Fischer

This dissertation studies three topics related to different types of network effects in financial economics. The first chapter, *Of Coupons and Cargo – International Debt, Production, and Trade*, quantifies the relationship between firms' supply chain networks and financing decisions. Most multinational corporations raise a significant amount of debt capital outside their home country. In contrast to prevailing evidence, access to deeper financial markets cannot explain this phenomenon in its entirety, as international debt issues carry higher spreads than securities concurrently issued domestically. Novel data on the universe of fixed income securities, subsidiary locations, and shipment-level trade flows from seventeen countries, is used to understand the drivers of international debt issuance. On the extensive margin firms raise debt in exactly those markets which play a key role in their supply chain (through subsidiaries, suppliers, or customers). Tests on the intensive margin indicate that firms adjust the face value of debt outstanding in a given country following exogenous changes in their operating exposure. These results are consistent with firms using international capital markets to hedge their exposure to fluctuations in exchange rates.

The second chapter, *Did You Catch the Game Last Night? – Peer Group Effects in Sell-Side Analyst Forecasts*, assesses the reaction of sell-side equity analysts to sentiment shocks, as well as how such non-financial information permeates through social networks. We identify a source of peer group influence that is plausibly orthogonal to information provision, yet nonetheless affects economic decision-making: the shock to an equity analyst of their undergraduate college football

team winning the NCAA Championship Game. We find that analysts' forecasts respond positively to their undergraduate school's football team winning the NCAA final. We then show that the shock of 'winning' spreads within an analyst's brokerage, positively influencing the forecasts of their colleagues. Brokerages where the degree of this diffusion is greater have lower female representation in their analyst teams, as well as lower ESG scores.

The third chapter, *Sharing is Caring? – Knowledge Diffusion in Researcher Networks*, focuses on the effects of social networks in innovation. Social interactions are at the core of many economic processes, including research and development. Yet their contribution to innovation is not well understood. A novel dataset on more than 19,000 economists linked to more than one million unique research projects and fifty million tweets (#EconTwitter) is used as laboratory to explore the relationship between different social interactions and research outcomes. Results suggest that interactions play a dominant role in the idea generation phase of research and a lesser one in the context of ongoing projects. They seem to matter little for completed research projects. More socially active scholars are more productive, as measured by the number of papers written, and their working papers are more visible (i.e., downloaded more frequently). A working paper being endorsed leads to an increase in downloads by 20%. However, indicative of a trade-off in spending their valuable time, these projects are less impactful based on citation measures.

Table of Contents

List of Figures

List of Tables

Acknowledgements

I am indebted to my advisers Tania Babina, Xavier Giroud, Jesse Schreger, Suresh Sundaresan, and Daniel Wolfenzon. Xavier has been an incredible mentor, making sure I did not lose faith, providing a light at the end of even the longest and darkest tunnel. He has been involved in my projects from early ideas to their execution, from big picture discussions all the way down to the nitty-gritty details. He found novelty in the questions I tried to answer and guided me in pitching them accordingly. Tania provided actionable advice on how to execute and present my ideas. By being my harshest critic, she made sure that my research was economically and econometrically sound. Daniel guided me diligently in my diverse research endeavors and offered his support whenever I needed it. Jesse, although a late addition to my committee, offered invaluable feedback and mentorship, which went far beyond my papers. Suresh, in his role as the PhD director, was the first person I ever spoke to at Columbia and who asked me if I could see myself moving to New York. In the early years he provided guidance in mastering the coursework, while later supporting me not only academically but also financially, for which I will be ever grateful.

The PhD journey would not have been the same without the friends I made along the way. Throughout the years I could always draw on academic as well as emotional support from my incredible peers Xiaobo Yu, Dhruv Singal, and Brian J. Lee. Dhruv, Xiaobo, and I have shared in each other's successes and found company in periods of misery. They are the closest thing to family I found in New York. Over the years, Brian listened to many a rant of mine and gave pep talks when needed. Thinking back to the early days in Uris Hall, I also owe a thank you to the final (unofficial) resident of cubicle 4QR, Jimmy (Ting-Yuan) Kuo, who, despite being an ocean away, has been an invaluable friend ever since. Outside the business school, I want to thank Eddie Shore, an amazing co-author and friend, as well as Daniel Markovits, a great roommate, loyal friend, and source of an outside perspective. Without any of them, I would not be who and where I am today.

Throughout my entire life my family has been a source of insatiable support. First and foremost, I am incredibly grateful to my parents, Sabine Knauer-Fischer and Markus Fischer, who have been supportive of all my endeavors and are the two people I look up to the most. I am indebted to my brother, Justus Benedikt Fischer, for always having my back and my grandmother, Erika Knauer, for a sympathetic ear whenever I needed it. My gratitude further extends to my other grandmother, Heike Fischer, my uncle, Matthias Fischer, as well as all remaining members of my family, who I know I can rely on unconditionally.

To those who matter most.

Chapter 1:

Of Coupons and Cargo

International Debt, Production, and Trade[1]

1.1 Introduction

International debt capital markets have grown rapidly over the past two decades, more than doubling in size since 2000. According to the Bank for International Settlements (2024), 36% of all debt securities are international as of Q4 2022, i.e., the lender is located outside the borrower's home market. Despite the incredible and prolonged importance of these financial markets, hard evidence on why a growing number of corporations choose to borrow abroad is scarce. A common explanation for firms borrowing internationally is that of requiring access to deeper financial markets. Under this point of view, firms located in countries with shallow financial markets eventually outgrow the capacity of domestic lenders, issuing in deeper foreign markets (e.g., U.S., Europe) instead. Despite more than five decades of research on public corporate debt markets, a lack of reliable data has been a major roadblock in quantifying differences between securities issued in different jurisdictions. In this paper, novel data on international fixed income instruments is combined with micro-data on imports and exports from 17 countries, as well as the locations of international subsidiaries to quantitatively approach the question why firms issue internationally.

Globalization has been one of the major phenomena shaping the twenty-first century. Considering the associated increase in international trade and relevance of integrated, cross-country

[1]This chapter is based on Fischer (2024). I am deeply indebted to my advisors Xavier Giroud, Jesse Schreger, and Daniel Wolfenzon. The paper benefited greatly from discussions with Kent Daniel, Amanda Dos Santos, Adam Rej, Dhruv Singal, Suresh Sundaresan, and Xiaobo Yu, as well as seminar participants at Columbia Business School, Barclays Quantitative Portfolio Strategy, and Cornerstone Research. An earlier version of this project was developed with Dhruv Singal during a joint research residency at Capital Fund Management (CFM) under the supervision of Adam Rej. I am immensely grateful to Suresh Sundaresan, who provided financial support.

Figure 1.1: International Debt and Trade

This figure depicts the time-series of quarterly changes in international debt issuance and international trade. Regressing quarterly changes in international debt outstanding on logarithm changes in trade volume, the coefficient on the latter has a value of 0.13 and is significant at the 1%-level.

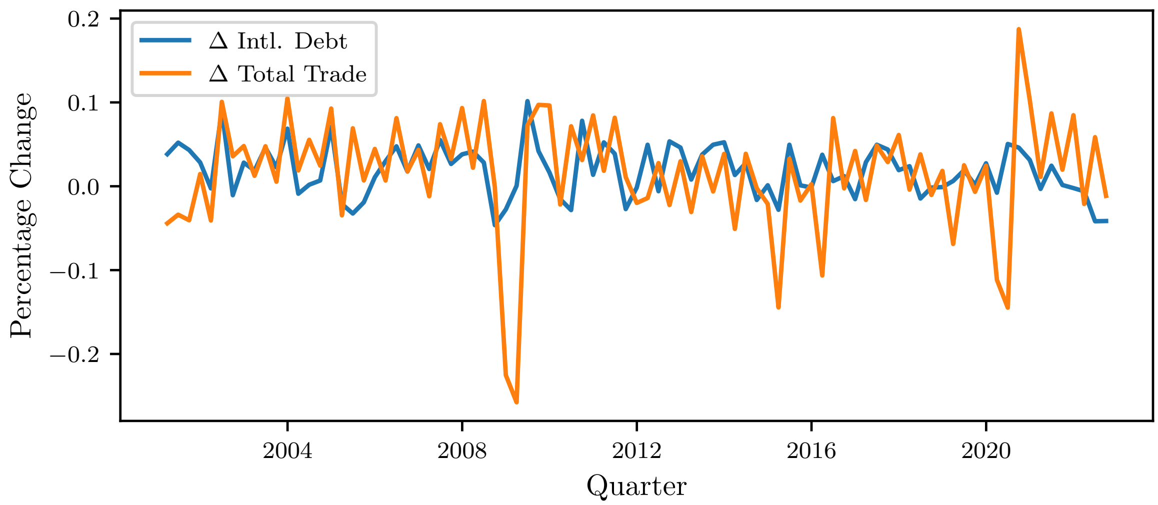

value chains, e.g., through the establishment of foreign subsidiaries, it is natural to hypothesize that these two phenomena are related. Figure 1.1 plots the quarterly growth rates of international debt and total trade. While the evolution of international debt is less volatile, it is difficult not to see co-movement between the two. A linear regression of international debt on trade reveals this, with a highly significant coefficient estimate of 0.13. In a friction-less world with a single currency, multinational corporations would choose to borrow in the market that affords them the lowest interest rates, channeling funds to wherever they are needed through their internal capital market. The real world differs from this benchmark in several ways: (1) firms are exposed to foreign currencies, (2) financial markets (domestic and international) are not friction-less, and (3) neither are internal capital markets (Dai et al., 2024; Shin and Stulz, 1998). Based on these high-level observations, this paper sheds light on the motivation behind international fixed income issuances and their relationship with the borrower's supply chain exposure in the time-series.

As a starting point, the anecdotal evidence of accessing deeper financial markets is put to the test. By comparing domestic to international issues by the same issuer with respect to their yields and spreads, the differences in borrowing cost across countries are quantified. Controlling for firm-year fixed effects, as well as security characteristics, a statistically significant difference in spreads of about 12 basis points (bps) between international and domestic instruments is documented. As international issues are seemingly more expensive, it is even more puzzling why large corporations would choose to tap foreign financial markets. To assess whether this result is driven by market timing, the analysis is repeated using yield quotes from the secondary market, which exhibit differences in spreads that are remarkably similar in sign and magnitude. By running the analysis separately for different countries/regions, it is assessed whether this effect is driven by a specific country. The effect is least pronounced for foreign firms issuing in the United States but prevails in all other markets. With existing research placing a strong emphasis on the U.S. capital market, it does not come as a surprise that a countervailing effect in other markets has remained overlooked. Comparing securities along other dimensions, it is documented that international issues are usually denominated in the currency of the country of issuance, suggesting their use in exchange rate management.

It seems intuitive that firms use international financial markets to fund foreign subsidiaries and hedge foreign exchange (FX) risk. This view is consistent with evidence in Keloharju and Niskanen (2001), who find that multinational firms are more likely to borrow in foreign currencies. In leveraging a wide range of novel data, this paper seeks to quantify this relationship at the firm-country-level.[2] To assess the relationship between international borrowing and corporate supply chains, data on imports/exports and the location of subsidiaries are utilized. As firms can serve foreign markets in different ways, e.g., through exports or setting up subsidiaries, it is important to consider these two dimensions in order to form a holistic view of a firm's supply chain.

The propensity to issue debt internationally is modelled as a linear function of a firm's (net) exposure to a foreign country, measured as excess exports and the number of subsidiaries, respectively. Consistent with the notion that firms use international capital markets to offset their trade exposure, corporations are most likely to borrow in exactly those countries with which they have

[2] Existing research is limited to the extensive margin, whereas this paper focuses on the intensive margin, i.e., the relationship between the exact destination of sales and the location of debt issuance.

large trade imbalances. This effect is driven by the long side, i.e., when exporting more to a country than importing from it. This is consistent with a natural hedging strategy, under which international fixed income securities are used to offset exposure to exchange rate fluctuations. A corporation that is net long in a foreign country can issue a claim in that jurisdiction, convert the proceeds to her home currency at the spot rate, and make interest and notional (re-)payments from her foreign cash flows. If this hypothesis indeed is a valid reason for international borrowing, the nominal value of debt outstanding in a foreign country should respond to changes in the trade exposure.

Taking the firm's decision whether to export or operate foreign subsidiaries (Horst, 1972) as given, the final set of results pertains to the response of the change in the nominal amount of debt outstanding to movements in the physical exposure. To establish a causal link between the firm's choice to issue internationally and the supply chain exposure, exogenous variation in the latter is needed. As the relationship is endogenous, an instrumental variable strategy is adopted. Using changes in tariffs to instrument for changes in trade is commonplace in the literature (Boehm et al., 2023). However, since the analysis in this paper is conducted at the firm-level, individual tariff lines must be aggregated for each corporation in the sample. Demonstrating that the resulting instrumental variable is valid, it is used to instrument for changes in trade exposure. Based on the point estimate, a one standard deviation shock to a firm's trade exposure results in a ten percent change in the amount of debt outstanding.

This paper provides evidence against a common explanation for cross-border borrowing, while proposing an alternative channel, which is upheld in the data. As access to deeper financial markets cannot rationalize the growth in international borrowing, corporations' desire to cheaply hedge FX exposures by entering natural hedges seems a more likely explanation.

Related Literature. This paper is closely related to three separate strands of literature in finance and economics. First, by considering the motivation behind raising corporate debt abroad, it adds to the long-standing corporate finance literature. Early research has focused on the optimal choice between debt and equity (Leland, 1994; Modigliani and Miller, 1958). Since then, scholars have shed light on how firms choose between different aspects of corporate debt, such as the number of creditors and their rights (Bolton and Scharfstein, 1996; Roberts and Sufi, 2009), seniority and maturity (Choi et al., 2018; Diamond, 1993), as well as the impact of monitoring (Park, 2000). Most recently Benmelech et al. (2022) focus on the choice between secured and unsecured debt, paying special attention to the trade-off between the two in the time-series. Closely related to this paper, Friedrich and Zator (2023) consider the impact of debt on corporate flexibility, following a shock to a firm's operations due to international tensions. Similarly, Allayannis et al. (2003) study the use of international debt by Asian firms around the Asian financial crisis, concluding that hedging motives did play a key role. Yet, to the best of my knowledge, there is no empirical evidence on the elasticity of debt with respect to a firm's exposure to a foreign country.

In considering the international aspect of corporate debt, this paper furthermore adds to the rapidly growing international finance literature. Benefiting from recent improvements in data availability across countries, the behavior of international financial markets and the firm's choice to partake in them are analyzed. Wang et al. (2020) focus on the difference in the (domestic) cost of capital between domestic and multinational corporates. International exposure as a source of risk and how firms are instrumental in transmitting it has been the focus of recent work by Richmond (2019) and Hassan et al. (2023). A growing literature explores the demand side of international debt markets. Using cross-country data on holdings of bonds, Coppola et al. (2021) redraw the map of international capital flows. Maggiori et al. (2020) document a home currency bias in the portfolio choice of fixed income investors. With respect to the supply side, early research by Keloharju and Niskanen (2001) on Finnish corporates suggests that international exposure is a key driver in foreign currency borrowing. Liao (2020) finds that corporations are more likely to issue in foreign markets if covered interest rate parity (CIP) is violated, making an arbitrage profit in the process. Closely related to this paper, Bruno and Shin (2023) consider the relationship between the markets for physical goods and those for financial products.

Lastly, by considering the relationship between a firm's financing and exporting decision, this paper is related to a large literature in international economics and trade. A long-standing question in this literature has been related to the effect of changes in tariffs on trade flows. As in Boehm et al. (2023), these analyses are usually done at the product-level. Their impact at the firm-level has been studied by Fitzgerald and Haller (2018). On the financial side, the international economics literature has predominantly been concerned with frictions in financial markets (e.g., Antràs and Caballero, 2009; Manova, 2012; Manova et al., 2015). A firm's choice how to deal with FX exposure is the topic of a famous case study by Antràs and Foley (2015). Firms exploiting financial frictions by engaging in a physical carry trade is corroborated by evidence in Liu et al. (2022).

The remainder of this paper is organized as follows: Section 1.2 describes the datasets used, the variable construction, and presents summary statistics. The empirical approach is outlined in Section 1.3. Section 1.4 presents the three sets of results. Ultimately, Section 1.5 concludes.

1.2 Data and Variables

To answer the question of what determines where corporates raise debt capital, a large dataset on corporate borrowing, subsidiaries, and the flow of goods is constructed. This section gives an overview of the data underpinning the analyses and how they are used to construct various measures of a firm's financial and physical exposure to a given country. The two standout sources of data used are those on debt securities and firm-level trade, providing a highly granular view of a firm's financing and operating activities.

1.2.1 Data Sources

Corporate Debt. All data related to corporate borrowing are obtained from Thomson/Refinitiv.[3] The dataset contains more than 8 million unique debt securities issued in over 120 different countries. See Figure 1.2a for an overview of the scale and scope of the data. For each security, the data contains information on the issuer, face value, coupon rate, coupon frequency, maturity, and

[3] Refinitiv started distributing its data products under the London Stock Exchange Group (LSEG) brand in early 2024.

any embedded options (put/call/conversion). After linking each issue to the global ultimate parent, those by financial companies are removed from the sample, leaving 918,738 issues.[4] See Table A.2 in the Appendix for at overview of the share of securities issued in a given country that could be linked to an ultimate parent in Compustat. The entire time-series of amounts outstanding is available for more than 99% of the sample. Whenever there are changes in this amount, the vendor records the reason for the change, the most common of which are repayment at maturity (75.67%) and the exercise of an embedded call option (12.95%). Table A.3 in the Appendix lists the prevalence of all corporate events leading to changes in the amount outstanding, as well as the distribution of the associated increases/decreases as a share of the original issuance amount. Where available, secondary market yields are collected as well. As most fixed income securities are illiquid, yields from different venues (i.e., exchanges, OTC) are combined to construct the most comprehensive time-series of trading yields possible. To control for time-varying differences in risk premia between tenors and changes in risk-free rates, government yields of the primary issuer in each currency are used to construct treasury yield curves.[5]

To obtain a more complete picture of the firms' debt composition, the fixed income data is supplemented with information on corporate loans and credit lines from DealScan. As these data predominantly capture syndicated lending, which almost exclusively takes place in a firm's home market, their exclusion does not materially affect the results. For each loan the borrower, amount, currency, maturity, and interest rate are recorded and treated the same way as the corresponding objects in the public debt sample.

Trade. Transaction-level records of goods moving across the borders of 17 countries are taken from S&P Panjiva. The dataset is compiled from Bills of Lading (BoL) and Air Waybills (AWB), documents essential in international trade and which are collected by a country's customs agents

[4]It is a known fact that financial corporations issue a majority of debt instruments. These are usually much shorter-term in nature than those issued by corporates, and thus naturally more numerous all else being equal.

[5]The extent to which government securities can proxy for risk-free rates has been subject of a long literature. The concerns raised therein predominantly apply to the 'specialness' of the U.S. dollar and the convenience yield associated with holding U.S. treasuries. While other measures have been proposed (e.g., van Binsbergen et al., 2022), they are not readily available for the majority of countries in the sample.

Figure 1.2: Geographic Coverage

This figure depicts the scale and scope of the fixed income and trade data used in this paper. Panel a visualizes the total nominal value of debt issued in each country over the sample period (2007-2022). Europe, the United States, and China are the largest debt capital markets in the sample based on total face value issued. The sample contains securities issued in most developed and emerging markets. Very few securities are issued in Africa and the Middle East. Panel b presents the coverage of the trade data, where the net exposure is defined as $\log_{10}(net\ imports)_c = \text{sgn}(imp_c - exp_c) \times \log_{10}(1 + |imp_c - exp_c|)$.

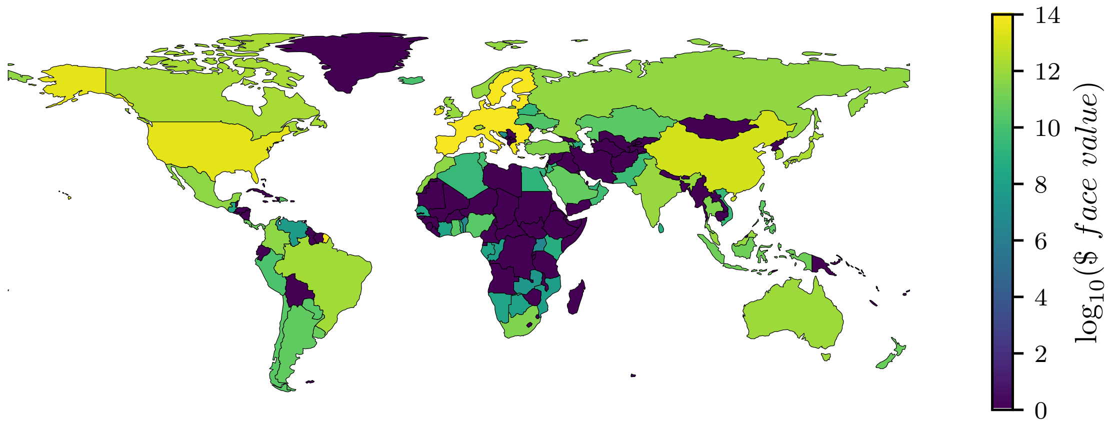

(a) Fixed Income Data

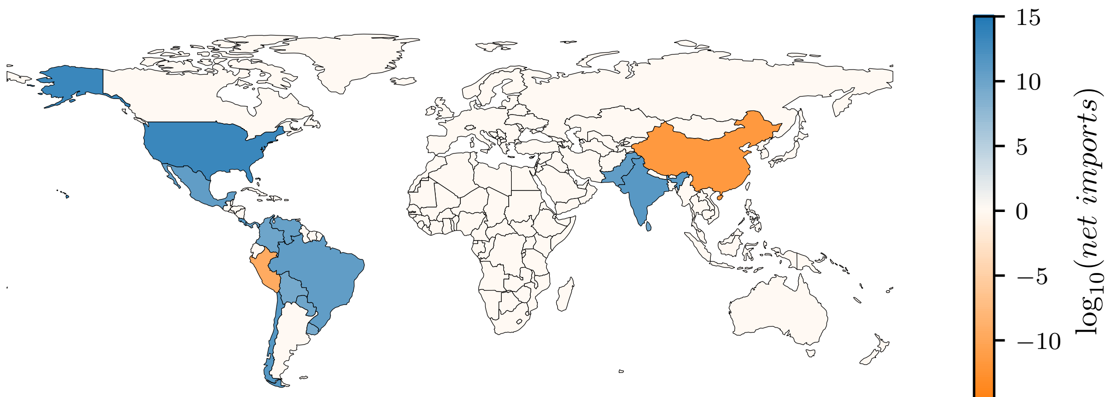

(b) Trade Data

to assess the value of goods and levy tariffs. While the exact resolution of the data differs between countries, the shipper, consignee, product, shipment quantity, weight, and value are generally available. The countries contained in the dataset are Bolivia (BO), Brazil (BR), Chile (CL), China (CN), Colombia (CO), Costa Rica (CR), Ecuador (EC), India (IN), Mexico (MX), Pakistan (PK), Panama (PA), Paraguay (PY), Peru (PE), Sri Lanka (LK), Uruguay (UY), the United States (US), and Venezuela (VE). For a visual representation of the countries in the data, as well as their classification as net importers/exporters, refer to Figure 1.2b. See Table A.1 and Figure A.1 in the appendix for an overview of the variables, transport modes, and years available for each country. The data contains more than 2.5 billion unique records and currently is the only commercially available source that provides information of this kind.[6] To give a sense of the origin of the variables contained in Panjiva, Figure 1.3 presents a picture of the form from which they are obtained in the case of U.S. imports. While this data has been used in answering questions pertaining to a single country, to the best of my knowledge, this is the first cross-country study using such data. For a detailed description of the U.S. data see Flaaen et al. (2023).[7]

Products in the trade data are identified through the harmonized system (HS) at the 6-digit-level, hereafter HS6. This level of granularity coincides with that at which most tariffs are set. While this classification allows for the cross-sectional comparison between product types, changes made to the system over time necessitate the translation between nomenclatures.[8] To link product types in the time-series, HS concordances are obtained from the UN Statistical Division (UNSD) and used to harmonize the data. To assess the cost of moving goods across borders, each shipment is linked to the applicable tariff. Annual, bilateral tariffs are collected from the UN's Trade Analysis Information System (TRAINS). Tariffs are generally set by the importer's country and specified at the level of the HS code. Both the harmonization of HS nomenclatures as well as the processing of tariffs follows Boehm et al. (2023).

[6] In an earlier version of this paper, similar data from IHS PIERS was used. The dataset was discontinued in 2022 after S&P acquired IHS Markit.

[7] One concern raised by the authors pertains to only maritime trade being observable. In order not to omit land-based trade between the United States and Mexico, such missing values are filled in using the Mexican data.

[8] The UN Statistical Division is responsible for the HS system and updates it every few years. Countries adopt these changes with varying lags, such that comparisons are only possible within a country-year.

Figure 1.3: CBP Form 1302

The data reported in S&P Panjiva are collected from bills of lading. In case of the United States, the vendor compiles the data by obtaining access to copies of CBP Form 1032 under the Freedom of Information Act of 1966 (FOIA). A blank version of such form is depicted below to give an idea of the information reported thereon.

DEPARTMENT OF HOMELAND SECURITY
U.S. Customs and Border Protection

OMB CONTROL NUMBER 1651-0001
EXPIRATION DATE: 11/30/2025

INWARD CARGO DECLARATION
19 CFR 4.7, 4.7a, 4.8, 4.33, 4.34, 4.38, 4.84, 4.85, 4.86, 4.91, 4.93, 4.99

Name of Vessel	Nationality of Ship	IMO No.	Voyage No.	Page No. of

Name of Master	Last Foreign Port Before U.S.	Port of Discharge	Date of Departure from Port of Loading	Time of Departure from Port of Loading (Zulu)

Shipper (SH) Consignee (CO) Notify address (NF)	Bill of Lading No.	Marks & Nos. (MN) Container Nos. (CN) Seal Nos. (SN)	No. & Kind of Packages Description of Goods Hazardous Materials (Must Provide UN Code)	Answer col. Gross Wt. (lb. or kg.)	Answer col. Measurement (per HTS)	First Port/Place Where Carrier Takes Possession of Cargo	Foreign Port Where Cargo is Laden on Board

Received by CBP (Signature): ____________________

Paperwork Reduction Act Statement: An agency may not conduct or sponsor an information collection and a person is not required to respond to this information unless it displays a current valid OMB control number and an expiration date. The control number for this collection is 1651-0001. The estimated average time to complete this application is 30 minutes. If you have any comments regarding the burden estimate you can write to U.S. Customs and Border Protection, Office of Regulations and Rulings, 799 9th Street, NW., Washington DC 20229.

CBP Form 1302 (12/22)

Subsidiaries. The locations, opening, and closing dates of domestic as well as foreign subsidiaries are obtained from the Bureau van Dijk's Orbis database. Combined with the parent linkages in Orbis, this allows for a holistic view of a firm's international operations. Each subsidiary is identified as one of 17 establishment types, allowing for the differentiation between operating and financing entities.[9] Doing so is of upmost importance, as the issuance of debt securities in foreign markets oftentimes necessitates the establishment of a local financing entity. All analyses will therefore exclude such establishments, unless explicitly specified otherwise. Furthermore, since the trade data only gives rise to a complete picture of a firm's exposure in the case of exporters, i.e., entities without a local presence, this data is used to exclude firms producing abroad when using trade as an exposure measure. For an overview of the types of entities and their relative frequencies in the

[9]Establishments are classified as one or more of the following: (B) Bank, (F) Financial company, (A) Insurance company, (C) Industrial company, (E) Fund, (J) Foundation/Research institute, (S) Public authority, (I) Individual/family, (M) Managers, (H) Self-owned, (P) Private equity, (Z) Listed, (D) Unnamed private shareholder, (L) Other unnamed shareholders, (Y) Hedge fund, (Q) Branch, and (W) Marine vessel.

data, see Table A.5 in the Appendix.

Miscellaneous. Balance sheet and income statement information is obtained from S&P Compustat Global, covering 60,924 unique entities between 2005 and 2022. To convert between currencies and control for expected changes in exchange rates, daily exchange and forward rates against the USD are taken from Refinitiv. The trade data is substantiated with information on ports, vessels, and their respective locations. Data on ports is obtained from the United Nations Economic Commission for Europe (UNCEE).[10] Details on vessels and their location are collected from the Automatic Identification System (AIS). The data is accessed through Refinitiv, with coverage of the largest vessels starting in 2010. This data is used to identify when shipments leave (arrive at) the port of lading (unlading).

1.2.2 Comparison to Alternative Data

Prior research on corporate borrowing in public debt markets has had a strong focus on the United States. A major reason for this bias is rooted in data availability. With granular data on international financial markets being scarce, the differences between securities placed in different countries have been left mostly unexplored until now. Covering the universe of corporate issuances, the data in this paper is, to a first order, a superset of that used in prior research. The natural point of comparison is Mergent's Fixed Income Securities Database (FISD), which is used in almost all research on the U.S. corporate fixed income market. There are about 300k securities that are present in Refinitiv as well as FISD. In Table A.4 in the Appendix, the relevant variables from FISD are regressed on their corresponding value as reported in Refinitiv. All coefficients are remarkably close to one, indicating that there are no systematic differences between the two datasets. Naturally, it is not possible to perform similar tests for securities issued in other countries. In untabulated tests, the data in Refinitiv is compared to a comparable dataset obtained from Bloomberg with

[10] A list of all ports and their UN/LOCODEs can be found on the website of the UNCEE.

equivalent results.[11]

How firms are affected by macro events due to their exposure to a foreign country has been of interest to academics for a long time. In financial economics, it is common to use data from income statements, broken down by region, and, more recently, information extracted from earnings calls to measure country exposure (Hassan et al., 2023). Either approach, however, conflates operational and financial exposure, making such data ill-suited for any analysis of the interplay between the two. Furthermore, they are generally based on geographic regions considered significant by the entity and therefore strongly influenced by the choice of organizational structure. While micro-data on physical trade has been used by trade economists for decades, it is only available through collaborating with government agencies and thereby naturally limited to firms located in a single country. At a higher level of aggregation, e.g., products or countries, bilateral trade data is more readily available but lacks the resolution required to focus on individual firms. The BoL data in S&P Panjiva is similar in spirit to the former, while providing significantly more breadth, as it is not limited to a single country. To alleviate concerns that the individual shipments cannot reconcile the total trade flows observable in official statistics, the logarithm of the total value of monthly bilateral trade in Panjiva is regressed on the same quantity taken from UN Comtrade. Figure A.2 in the appendix presents the point estimates and confidence intervals obtained from these regressions for all countries in the sample. Since for some countries only a subset of all shipments is collected, e.g., only maritime trade in the case of the U.S. and Brazil (see Table A.1), naturally $\widehat{\beta} < 1$. That being said, the granular shipping data still captures more than 75% of total trade for all countries in the sample.

1.2.3 Public Debt Markets

Corporate debt comes in many forms and sizes. For the purpose of this paper, the major differentiator between various kinds of securities lies in their maturity. The fixed income data covers

[11] Data from Bloomberg's fixed income database was used in an earlier version of this paper. Before losing access to the Bloomberg data, extensive comparisons between it and Refinitiv were undertaken. The latter has a greater coverage and coincides almost perfectly with the former in the overlapping part of the sample.

instruments with maturities ranging from a few days to many decades. Securities with maturities of less than one year are classified as **commercial paper (CP)**. These instruments do not pay interest but are issued at a discount, similar to zero-coupon bonds (ZCB). Fixed income instruments with maturities between one and ten years are referred to as **medium-term notes (MTN)**. Depending on their maturity, these securities can be either discount notes or pay coupons in regular intervals. See Figure A.3 in the Appendix for an overview of the share of discount securities for different maturities. **Corporate bonds (CB)** sit at the upper end of the maturity spectrum, with more than ten years to maturity at issuance. All of these securities pay coupons, whose frequency ranges from annual to monthly.

As is the case in equity markets, there exists a meaningful difference between the primary and secondary market for corporate debt issues. As it pertains to the analyses in this paper, **primary market** yields closely reflect the actual cost of borrowing at the time of issuance. Differences between the offering yield and the cost of borrowing are attributable to a non-zero spread charged by the underwriter. Absent data on these spreads, primary market yields are a natural lower bound of borrowing costs. While being a comparatively clean measure of the cost of debt, these yields are only observable at issuance and therefore do not give rise to a meaningful time series. In contrast, **secondary market** yields are derived from the prices at which a firm's securities are trading at any given point in time. They reflect the current assessment of a firm's risk and therefore can be interpreted as a proxy for the marginal cost of borrowing for a specific maturity. However, they are only available for a selected subset of securities, hence all analyses will be performed on primary as well as secondary market yields.

1.2.4 Corporate Internationalization

While this paper mostly abstracts away from a firm's decision how to serve international customers, it is important to have a broad understanding of corporate internationalization strategies. Whenever a corporation chooses to serve customers in a (foreign) country c, it has to decide on the means of doing so. Figure 1.4 outlines a possible ordering of internationalization strategies based

Figure 1.4: Corporate Internationalization

Corporates have a wide range of possible market entry strategies at their disposal. They differ in the degree of control the firm has over operations and the risks assumed. Adapted from Kreikebaum (2003).

on Kreikebaum (2003). For the purpose of assessing the relationship between international operations and corporate financing decisions, the two extreme ends of the scale are most meaningful. **Exporters** ship their goods to customers located in foreign countries. In doing so, they must fund their (short-term) working capital needs. This leads to a natural incentive to use foreign short-term debt markets as a means of funding working capital tied up in the shipment process and manage exchange rate risk. Entities for which this applies will be analyzed by restricting the sample to corporations without local subsidiaries. On the other extreme, firms may establish **subsidiaries** in countries which represent a substantial part of their market. The automotive sector can be seen as a canonical example of this. Over the last decades, many car manufacturers have set up factories in the U.S. and China in order to serve these markets locally and avoid the significant shipping costs and tariffs associated with moving large goods internationally.

1.2.5 Variable Construction

Borrowing Cost. For 539,588 (58.73%) of the corporate debt instruments in the sample, all variables needed to compute the yield to maturity at issuance are available. The yield to maturity y is defined as the rate of return at which the present value of future coupon payments and the repayment of the face value equal the market price.

$$P = \sum_t \frac{C}{(1+y)^t} + \frac{F}{(1+y)^T}$$

y furthermore corresponds to the return to the investor if the security is held to maturity but might differ from this value in the case of trading. With the exception of unobservable proceeds charged by the underwriter, it should closely correspond to the borrowing cost faced by the issuer. In addition to the issuance yield y, monthly trading yields are observable for 223,354 unique securities. With 6.55 million monthly observations in total, there are on average 29.33 monthly yields for each security.

Since yield curves for different currencies can differ widely, it is paramount to make comparisons solely based on the spread over the corresponding risk-free rate. The government curve is derived from *on-the-run* government securities and interpolated for all intermediate tenors. In order to obtain a more reliable measure of the spread in the primary as well as secondary market, the procedure proposed by Nelson and Siegel (1987) and refined by Svensson (1994) is used to compute the maturity-matched risk-free rate in the same currency. See Section A.5 in the Appendix for a more detailed overview of how this method is applied to construct yield curves.

Financial Exposure. To turn the outstanding amount of debt securities into a firm-country measure of financial exposure, each instrument must be linked to a specific country. If available, the country assigned by Refinitiv is used. For those securities without a country assigned to them, the approach outlined in Petre (2009) and followed by the BIS is applied. In a nutshell, the first two characters of the ISIN are used to determine the market/country a security has been issued in. Where this is not feasible, the location of the issuer, guarantor, and primary exchange (for exchange listed bonds) are used instead. With all instruments assigned to an issuer as well as a primary market, the exposure of firm f to country c in period t is computed as the total nominal value of debt outstanding.

$$financial^{\$}_{c,f,t} = \frac{1}{|\mathcal{T}_t|} \sum_{b \in \mathcal{B}_{c,f}} \sum_{d \in \mathcal{T}_t} outstanding_{b,d} \times fx^{\$}_{u,d}$$

Let $\mathcal{B}_{c,f}$ denote the set of securities issued by firm f in country c, and $\mathcal{T}_t$ the set of days in year t. $outstanding_{b,d}$ is the outstanding balance of security b on day d. As firms may, and do, repay

Figure 1.5: International Trade Flows

This figure describes two ways in which a firm's exposure to a foreign country can be described, based on the example for shipments from/to Brazil. Panel a represents the case in which a foreign company exports to/imports from Brazil without operating any subsidiaries there. In this case, the trade records correspond to the total exposure of the international company to Brazil, regardless of the location of their subsidiaries responsible for the shipments. As depicted in panel b, for all domestic companies (i.e., those without establishments in foreign countries) the shipments reported in the Brazilian data correspond to their total exposure.

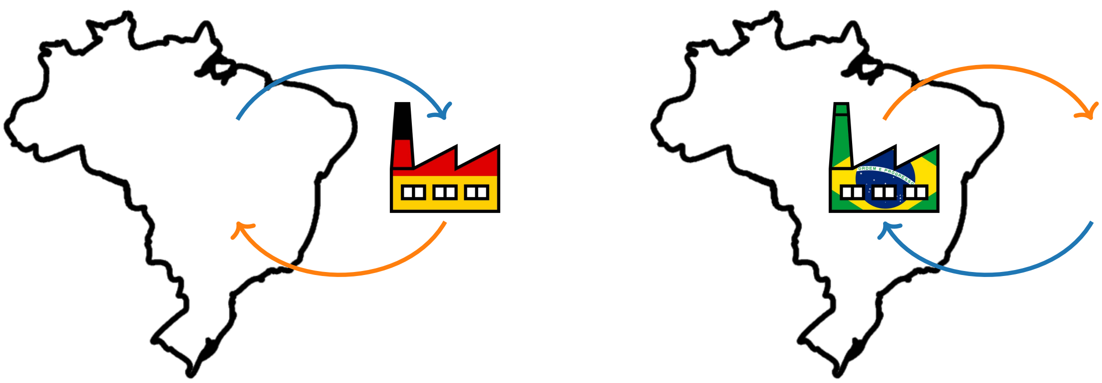

(a) Multinational Firms **(b)** Domestic Firms

issues early (see Table A.3), the outstanding amount changes over time. This value is converted to USD at the daily spot exchange rate $fx^{\$}_{u,d}$ for currency u. As the trade data contains the shipment values in USD, this alleviates concerns that results are driven by exchange rate fluctuations and makes magnitudes comparable.

Trade Exposure. After linking all shipments in the trade data to the corporations involved, it must be decided how to measure a firm's trade exposure to a given country. The granularity of the shipping data allows for the construction of different kinds of exposure measures. While research in trade economics has historically been concerned with the flow of goods between countries, this paper considers trade flows between firms. Therefore, each shipment has four dimensions: (1) the shipper, (2) the country of origin, (3) the recipient/consignee, and (4) the destination country. Figure 1.5 presents two natural ways of approaching the data.

Firms can be classified as either domestic or multinational, depending on the existence and location of any subsidiaries. If a firm does not have any foreign subsidiaries, it can be considered domestic, as in Figure 1.5b, with all shipments to/from foreign countries naturally corresponding to the physical exposure. However, as these firms are on average smaller, they are also less likely to borrow from public markets. By looking at international debt markets, it therefore is natural to consider multinational corporations (MNCs) instead. As depicted in Figure 1.5a, if a firm does not have any subsidiaries in any of the 17 countries for which their shipments are observable, the import and export data of that country is a clean measure of the firm's exposure. The sample therefore is restricted to the latter.

Having established a notion of *country exposure*, the next question pertains to the quantification thereof. Naturally, trade can take the form of imports or exports, which cleanly map to the financial concepts of long and short exposures. Let $\mathcal{C}$ and $\mathcal{F}$ be the set of countries and firms, respectively. The long end corresponds to the creation of receivables in the foreign country (and currency) and can be computed as the total value of shipments from any subsidiary of firm f in country d to a consignee g in country c.

$$exposure^{long}_{c,f,t} = \sum_{d \in \mathcal{C}} \sum_{g \in \mathcal{F}} shipment^{\$}_{d \to c, f \to g, t}$$

Similarly, a short exposure (i.e., payable) corresponds to the total value of shipments from any corporation g in country c to all subsidiaries of firm f in countries $d \in \mathcal{C}$.

$$exposure^{short}_{c,f,t} = \sum_{d \in \mathcal{C}} \sum_{g \in \mathcal{F}} shipment^{\$}_{c \to d, g \to f, t}$$

From these definitions it is immediately obvious that the exposure of the focal firm f to country c does not depend on the home country. Substituting $\mathcal{F} \setminus \{f\}$ for $\mathcal{F}$ in the expressions above, one obtains a measure of shipments to/from unaffiliated corporations. If one defines $\mathcal{F} = \{f\}$, the measures reduce to the total value of intra-company shipments. As transfer pricing between affiliated entities has been an area of great interest, both from academics as well as practitioners (Baersch et al., 2018), it is important to be able to isolate shipments for which these concerns might play a significant role. Since long and short exposures can offset one another, the net exposure is defined as the difference between them.

$$\ln\left(net\ exposure\right)_{i,c,t} = \mathrm{sgn}(imp_{i,c,t} - exp_{i,c,t}) \times \ln\left(1 + |imp_{i,c,t} - exp_{i,c,t}|\right)$$

This measure of exposure has a natural inflection point at a value of zero. It will therefore only be included in regression when interacted with a dummy of whether the exposure is positive or negative.

Subsidiaries. When choosing to enter a foreign market, corporations have a range of market entry strategies at their disposal (Helpman et al., 2004; Kreikebaum, 2003). Shipping goods to the country in which customers are located is the mode of entry with the least control and long-term financing needs (see Figure 1.4). As ties to foreign countries mature, the establishment of local subsidiaries, which produce goods to serve the local market directly, becomes more advantageous. The existing literature on international corporate borrowing predominantly compares the choice of borrowing between firms with and without substantial international presence (Keloharju and Niskanen, 2001). To assess the effects the the firm-country-level, more granular data is needed. Using detailed establishment data from Orbis, the physical presence of a firm in a foreign country can be assessed. A major benefit of the granularity of this dataset lies in the differentiation between several types of establishments. As the issuance of international debt instruments might require the establishment of a local financing company, these entities are considered separately from manufacturing establishments. This differentiation is a simple one to make in the data, allowing for the construction of a more meaningful measure of exposure as the number of producing subsidiaries in each country.

1.2.6 Summary Statistics

The analyses in this paper rely on two separate datasets, (1) issue-level fixed income data from the primary/secondary market, and (2) firm-country measures of financial and trade exposure in the time-series. This section therefore in turn presents summary statistics for the two types of data.

Debt Securities. Table 1.1 gives an overview of the fixed income sample and the variables therein, split by whether the issues are domestic or international. As security characteristics are determined jointly, the differences in means between the two types of issues in the sample should be interpreted with caution. Nevertheless, they can help in understanding reasons for why firms decide to issue debt in foreign markets. Panel A presents summary statistics for the primary market data, i.e., characteristics at the time of issuance.

At issuance, international debt securities are smaller and have significantly shorter tenors than those issued domestically. Contradicting anecdotal evidence, both yields and spreads are significantly higher for international issues, however, this result might in part be driven by differences in maturities and issuers, necessitating a more nuanced assessment of the differences in spreads.

An important takeaway from Table 1.1 pertains to the currencies of the securities. While it is impossible to know if proceeds are converted to foreign currencies after issuance, considering the propensity of international issues being denominated in the market's local currency is nonetheless important. While the difference in this dummy variable between the two sub-samples is statistically significant, the mean of both variables is close to one, indicating that debt securities are generally issued in the local currency of the financial market they are placed in. This effect is not driven by securities issued in the U.S. but pervasive across countries.

In Panel B, which contains the summary statistics for the secondary market data, most observations related to yields and spreads remain largely unchanged. Some of the differences in characteristics change sign when moving from the primary to the secondary market data. Comparing Panel A to Panel B, it can be seen that traded securities are significantly larger than the average bond, while also having much longer tenors. They furthermore carry lower yields and spreads,

Table 1.1: Summary Statistics – Fixed Income Sample

This table presents summary statistics of the fixed income sample, split into domestic and international securities. Issues are classified as domestic if they are placed in the issuer's country of incorporation. International securities are those issued in foreign financial markets.

	International		Domestic		Δ
	Mean	S.D.	Mean	S.D.	
Panel A: Primary Market	*(N=294,595)*		*(N=245,037)*		
$\log_{10}(outstanding)$	5.98	0.72	6.67	0.99	-0.69***
Tenor (months)	24.20	37.42	41.45	42.55	-17.25***
CP	0.48	0.50	0.23	0.42	0.25***
MTN	0.49	0.50	0.73	0.45	-0.24***
Bond	0.03	0.17	0.04	0.20	-0.01***
Senior	0.55	0.50	0.36	0.48	0.18***
Secured	0.01	0.08	0.03	0.17	-0.02***
Local currency	1.00	0.00	1.00	0.00	0.00***
Yield (bps)	759.55	514.69	577.03	475.21	182.52***
Spread (bps)	675.72	524.98	430.73	500.60	244.99***
Panel B: Secondary Market	*(N=2,417,178)*		*(N=563,701)*		
$\log_{10}(outstanding)$	8.02	0.90	7.86	0.75	0.15***
Tenor (months)	102.54	86.13	91.85	75.03	10.69***
CP	0.02	0.14	0.01	0.09	0.01***
MTN	0.71	0.45	0.78	0.41	-0.07***
Bond	0.27	0.44	0.21	0.41	0.06***
Senior	0.75	0.43	0.77	0.42	-0.02***
Secured	0.18	0.39	0.09	0.28	0.09***
Local currency	0.97	0.17	1.00	0.07	-0.02***
Yield (bps)	407.99	417.96	288.97	299.08	119.02***
Spread (bps)	203.46	394.21	145.90	265.31	57.56***

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

which, at least in part, is attributable to their increased liquidity.

Sample Composition. Due to the way the data in this paper is constructed, firms entering the sample are located in a wide range of countries. Table 1.2 presents an overview of the ten countries with the greatest number of firms represented in the matched data. For a firm to enter the sample it must have financial data available in Compustat and at least one security outstanding over the sample period (2007-2022). Firms located in the U.S., Korea, Japan, and China are the most prevalent in the data. The EU, when taken together, would be the fourth largest contributor to the sample by

Table 1.2: Summary Statistics – Sample Composition

This table presents the distribution of firms in the sample. Each firm is assigned a country of residence based on the location of their headquarters. Only firms for which financial are available in Compustat and which have at least one fixed income security outstanding over the 2007-2020 period enter the sample. When taken together, there are 1,166 located in the European Union, which issued a combined 190,894 debt securities.

Country	**#Firms**	**#Issues**
US	2,825	378,971
KR	1,621	21,672
JP	1,316	63,764
CN	1,148	12,014
TW	910	6,608
IN	622	27,497
CA	407	41,094
GB	255	62,207
MY	225	5,609
BR	211	4,163
FR	186	67,768
TH	169	2,061
DE	168	64,278
HK	156	1,851
Other	2,169	159,181
Total	12,388	918,738

number of entities and second largest, behind only the U.S., in the number of securities issued.

Country Exposure. Table 1.3 presents the country exposure variables constructed from the fixed income and trade data. As the firms under consideration are MNCs (as in Figure 1.5a), there naturally are at most 17 observations for each firm-year, one for each country represented in S&P Panjiva. As explained in the previous section, in addition to aggregating the USD value of imports/exports related to any given firm, the value of shipments within the boundary of the firm as well as those across them are computed separately.

Panel A contains the summary statistics for the trade exposure, which has been constructed using the trade data. Due to data limitation, it is not possible to compute both the long and short exposure for all firm-country-years. The average firm has a net long position that is increasing over time. Recall that the data is constructed such that only firms without subsidiaries in each

Table 1.3: Summary Statistics – Country Exposure

This table presents summary statistics for the sample of multinational corporations (MNCs). For each MNCs, there are at most 17 firm-year observations, one for each of the countries in the trade data. Countries in which the MNC has producing subsidiaries are removed from the sample.

	N	Mean	SD	p25	Median	p75
Trade						
$\ln(exposure^{long})$	418,901	9.40	5.97	4.83	11.12	13.79
$\Delta \ln(exposure^{long})$	317,811	0.14	3.60	-0.54	0.00	0.57
$\ln(exposure^{long,extra})$	418,901	9.16	5.95	0.14	10.98	13.61
$\Delta \ln(exposure^{long,extra})$	317,811	0.13	3.65	-0.55	0.00	0.57
$\ln(exposure^{short})$	295,499	4.39	6.11	0.00	0.00	10.59
$\Delta \ln(exposure^{short})$	217,327	-0.04	3.55	0.00	0.00	0.00
$\ln(exposure^{short,extra})$	295,499	4.25	6.00	0.00	0.00	10.38
$\Delta \ln(exposure^{short,extra})$	217,327	-0.04	3.57	0.00	0.00	0.00
Debt						
$\ln(exposure^{fin})$	418,901	0.30	2.48	0.00	0.00	0.00
$\Delta \ln(exposure^{fin})$	418,901	0.02	0.85	0.00	0.00	0.00
Subsidiaries						
$\ln(subs_{tot})$	1,826,366	1.23	0.85	0.69	0.69	1.39
$\ln(subs_{cor})$	1,826,366	1.15	0.85	0.69	0.69	1.39
$\ln(subs_{fin})$	1,826,366	0.17	0.45	0.00	0.00	0.00
Controls						
$\ln(assets)$	337,918	6.70	2.19	5.16	6.57	8.15
$\ln(debt)$	337,219	5.93	2.37	4.20	5.79	7.56
$\ln(cash)$	328,305	4.27	2.18	2.73	4.26	5.75
$\ln(capex)$	334,516	3.38	2.09	1.75	3.22	4.79
$\#subsidiaries$	192,273	9.75	56.85	0.00	1.00	3.00

country enter the sample. Moving on to the financial exposure in Panel B, it becomes obvious that only very few firms issue debt in the sample countries. The results therefore will naturally be driven by the largest firms, which is consistent with the prior that only these firms are able to routinely place securities in foreign capital markets. Like the average increase in trade exposure, the financial exposure increases over the sample period. While part of this could be related to changes in exchange rates, the U.S. Dollar Index (USDX), which proxies for the strength of the USD vis-à-vis a basket of currencies, was increasing over the period. Ceteris paribus, the USD value of assets denominated in foreign currencies should therefore have decreased. Lastly, Panel

C presents control variables at the firm-level, obtained from Compustat and converted to USD.

1.3 Empirical Strategy

To answer the question why firms choose to issue debt in foreign markets, three sets of statistical tests are undertaken. First, the anecdotal evidence that corporations do so to access deeper financial markets is assessed by comparing characteristics and yields of domestic and international securities. Thereafter, the extensive margin, i.e., the decision where to issue, will be analyzed at the firm-level. Lastly, a causal link between a firm's trade and financial exposure is established. As this relationship is endogenous, an instrumental variable strategy is outlined and applied.

1.3.1 Security Characteristics

As is evident from Table 1.1, securities issued in domestic and foreign markets are very different from one another. However, as security characteristics are jointly determined, one should not take the differences in means at face value. While individual instrument characteristics are endogenous, the equilibrium pricing of debt claims allows for a cleaner comparison between domestic and international issues, controlling for any other observables. To this end, the model in Benmelech et al. (2022) is adapted, augmenting the regression with a dummy variable that takes on a value of one for international issues and zero otherwise. Leveraging the granularity of the debt data, this comparison can be made based on either the primary market offering yield or the trading yield in the secondary market. Using differences in trading spreads, the implied difference in cost of debt capital can be quantified outside issuance events. In either setting the main regression specification takes the following form.

$$spread_{b,c,u,t,\tau} = \beta \times \mathbb{1}\{foreign\}_{b,c,t} + \Gamma X + \delta_{c,t} + \delta_{u,t} + \delta_{f,t} + \varepsilon_{b,c,u,t,\tau} \tag{1.1}$$

In the equation above, $spread_{b,c,u,t,\tau}$ is the spread between the (offering or trading) yield of security b over the corresponding risk-free rate.[12] The coefficient of interest is β, which multiplies a dummy variable that captures whether security b has been placed in the issuer's home country or in a foreign market. X contains controls for differences in observables between domestic and international issues. To further account for time-varying differences in firm characteristics a firm-year fixed effect $\delta_{f,t}$ is included. Changes in the overall composition of the capital market in country c and unique characteristics of different currencies u, are controlled for through the inclusion of country-year ($\delta_{c,t}$) and currency-year ($\delta_{u,t}$) fixed effects.

1.3.2 Extensive Margin

The existing literature has focused on the extensive margin, i.e., the high-level determinants in a firm's choice to issue debt securities in a foreign currency or a foreign market. To seamlessly connect the subsequent results to the existing literature, extensive margin analyses are conducted. In contrast to the prior literature, the relationship between physical and financial exposure is not conducted at the firm, but at the firm-country-level. Following Keloharju and Niskanen (2001), a dummy for whether a firm issues securities internationally is regressed on firm-level variables.

$$\mathbb{1}\{foreign\_bond\}_{f,c,t} = \alpha + \beta \times exposure_{f,c,t} + \Gamma X + \delta_{f,t} + \varepsilon_{f,c,t} \tag{1.2}$$

While the inclusion of firm-year fixed effects ensures that the estimates are not driven by firm-level unobservables, the result does not allow for a causal interpretation, as the choice to engage with foreign countries is naturally endogenous. These are the same concerns that have plagued the existing literature. Nevertheless, in performing this analysis at a more granular, firm-country-year-level, the bar for confounding effects to impact the results is higher.

[12] Risk-free rates are obtained from the most liquid government securities denominated in the same currency.

1.3.3 Intensive Margin

To better understand which corporations are most likely to issue debt securities in foreign markets, several models are estimated. The general regression specification is similar to that in Bruno and Shin (2023).

$$debt_{f,c,t} = \beta \times trade_{f,c,t} + \Gamma X + \delta_{f,c} + \varepsilon_{i,c,t} \tag{1.3}$$

As firms jointly choose their trade and debt exposures, an instrumental variable strategy is employed. Borrowing from the trade literature, changes in tariffs, which are plausibly exogenous to the individual firms in the sample, are used to instrument for trade exposure. To quantify the response of the total value of debt outstanding in country c, as the trade exposure changes, equation (1.3) is run in first differences.

$$\Delta exposure^{fin}_{f,c,t} = \beta \times \Delta exposure^{trade}_{f,c,t} + \Gamma X + \delta_{f,c} + \delta_t + \varepsilon_{i,c,t}. \tag{1.4}$$

Any estimation of β is subject to the standard endogeneity concerns – there are numerous common channels affecting both trade exposure and debt outstanding. To estimate the causal effect of changes in trade exposure on changes in financial exposure, an instrumental variable (IV) strategy is employed. In the follow section, the construction of the instrument is described in detail.

Instrument Construction

The construction of the instrument follows the theoretical foundation of international trade laid out by Melitz (2003). In their peper, trade between two entities is modeled as a function of a series of inputs, an important one of which is the cost of moving goods between countries. Costs are a composite of tariffs, shipping costs, insurance, and iceberg costs. Of those, tariffs have been shown to have a sizable impact on trade volumes in the aggregate (Boehm et al., 2023). It therefore is only natural to look towards tariffs in search of an instrument for trade exposure.

While the use of tariffs as an instrument for trade is not unique to this paper, the prior literature does so for individual products. As the analysis in this paper is at the firm-country-level, the item-specific tariffs must be aggregated. Following Boehm et al. (2023), the instrument for firm f's shipments between countries d and c in year t is defined as the year-on-year log-difference between in effective tariffs.

$$z_{f,t}^{d\to c} = \ln\left(1 + \widehat{\tau}^{eff}_{f,d\to c,t,t-1}\right) - \ln\left(1 + \widehat{\tau}^{eff}_{f,d\to c,t-1,t-1}\right)$$

The effective tariff $\widehat{\tau}^{eff}_{f,d\to c,t_\tau,t_w}$ is the weighted tariff applicable to goods moved from country d to country c by firm f, using tariff rates from period t_τ and weights from period t_w. This aggregation gives rise to an instrument as described in Borusyak et al. (2021), where identification comes from the exogeneity of the tariff changes. Put mathematically, the tariffs for products with HS code h are aggregated as follows.

$$\widehat{\tau}^{eff}_{f,d\to c,t_\tau,t_w} = \sum_h \tau_{d\to c,h,t_\tau} \times w_{f,c,h,t_w}.$$

w are the firm-specific weights used to aggregate the product tariffs τ_h to the firm-country-year-level. The weights must be chosen such that they are exogenous to the firm. Just like BLP-instruments (Berry et al., 1995), they are computed based on the imports (or exports) of all firms $f' \in \{F \setminus f\}$.

$$w_{f,c,h,t} = \frac{\sum_j v^{\$}_{j,c,h,t}}{\sum_j \sum_h v^{\$}_{j,c,h,t}}$$

In the expression above, $v^{\$}_{j,c,h,t}$ is the USD value of shipments by firm f into country c, in a product group h in period t. To ensure that the choice of weights does not introduce any forward-looking information, the weights are computed based on $t-1$ information and applied to compute effective tariffs for $t-1$ as well as t (e.g., Backus, 2020).

For the estimates obtained from equation (1.4) to allow for a causal interpretation, the usual requirements for instrumental variables must be met, (1) exogeneity of the instrument, (2) a significant first stage, and (3) the exclusion restriction. The case for exogeneity will be made using the construction of the instrument and the market structure. The first stage will be formally tested in the following section, while a case will be made while the latter is unlikely to be violated.

Strict Exogeneity

Exogeneity of the instrument could be violated if there is an endogenous, contemporaneous response of tariffs to trade or financial outcomes at the firm-country-year-level. This effect could work through either the tariffs τ_h or the weights w. Violation of strict exogeneity through the weights is ruled out by construction. While firms may endogenously respond to (anticipated) changes in tariffs, the weights used in aggregating tariff changes are kept constant, making any such concern ill-founded.

The exogeneity of the instrument may not be given if firms lobby for changes in tariffs. While it is difficult to rule out such an effect entirely, there is reason to believe that this concern is not of a first order in this setting. Firms in whose case this might be concerning are those that are responsible for a significant share of trade. Examining the share of total imports in each product category attributable to a single firm, see Figure 1.6, it becomes clear that individual firms in general are not large enough for this to matter. While lobbying might also originate from small firms, it is significantly less likely that these will have a meaningful impact on how tariffs are being set.

In addition to firms in the sample being small relative to the total size of imports and exports, there is further reason to believe that the exogeneity of the instrument is unlikely to be violated. Since the models to be estimated using this IV design include fixed effects, the exogeneity condition must only be satisfied conditionally, i.e., after controlling for all fixed effects that are included in equation (1.4). As these remove unobservable variation at the firm-country and year-level, any such effects would have to be time-varying within a firm-country pair. While this is still possible, it is practically infeasible for such a confounding endogenous time-varying factor affecting tariffs and trade flows to be salient in this setting.

Figure 1.6: Share of Shipments

This figure depicts the share of all imports/exports in a product category attributable to a single firm. From top to bottom, the measures depicted are the value of shipments (in USD), weight (in kg), and volume (in TEUs). Note that the y-axis is logarithmic. It can be concluded that there are effectively no firms responsible for more than single-digit percentages of total trade volume in a product group. This alleviates concerns that lobbying by firms is prevalent at scale in the data.

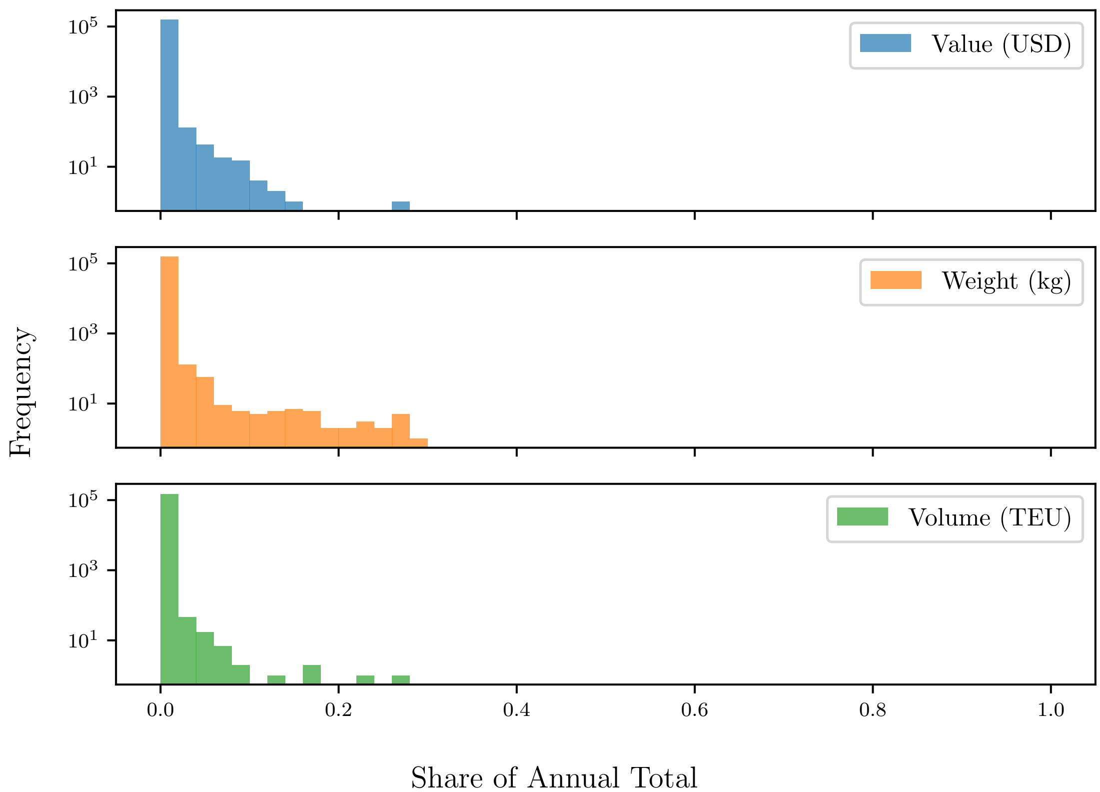

First Stage

The second requirement the instrument must satisfy is the relevance condition, i.e., it must have a significant effect on the endogenous variable. Fortunately, this requirement is easily testable by regressing the endogenous variable on the instrument. The first stage therefore takes the form below.

$$\Delta exposure^{trade}_{f,c,t} = \gamma \times z^{c,c'}_{f,t} + \Gamma X + \delta_{f,c} + \delta_t + \varepsilon_{f,c,t}.$$

The regression estimates presented in Table 1.4 are to be regarded as evidence that the instrumental variable strategy delivers a significant first stage. Naturally, the first stage must be significant for the sample the second stage is estimated on. As expected, and consistent with the prior literature, tariff increases lead to a reduction in shipments to a foreign country. In column 1, the dependent variable is computed over all shipments, regardless of whether they cross the firm boundary or not. Columns 2 and 3 are run on the exposure measures, considering only extra-firm and intra-firm shipments, respectively. While the sign of the effect is the same across specifications, the relationship is significantly less strong for intra-company shipments. That is not surprising, considering a firm's ability to use transfer pricing effectively, reduce their tariff burden by adjusting prices accordingly. Note as well that the adjusted R-squared in column 3 is negative, highlighting the that the majority is driven by the fixed effects, with little explanatory power coming from the instrument itself. It is for this reason that all analyses in this paper will be restricted to exposure measures based on all shipments and extra-firm shipments only.

Exclusion Restriction

Finally, the instrument must satisfy the exclusion restriction, i.e., the only way in which the change in weighted tariffs is allowed to move the financial exposure is through its effect on trade. As before, due to the inclusion of fixed effects in equation (1.4), this condition need only be satisfied after controlling for all fixed effects. While there are many ways in which the unconditional exclusion restriction could be violated, it is significantly more difficult to argue for a channel other than the one through trade once controlling for firm-country effects.

Any such channels would have to be time-varying within a firm-country firm, as they would otherwise be absorbed by the included fixed effects. That is, if the tariff changes are to affect financing changes via channels other than trade flow changes, these channels must be systematically present over time for the same firm-country pair. Thus, the exclusion restriction can be reduced to an absence of such effects, i.e., changes in tariffs either affect the financial markets due to their impact on trade flows, or due to systematic forces which are constant within a firm-country pair.

Table 1.4: Tariff Instrument – First Stage

This table reports the coefficient estimates for the first stage, using the change in effective tariffs as an instrument for the change in trade exposure. The regression estimates presented here are to be regarded as evidence that the instrumental variable strategy delivers a significant first stage. However, since the relevance criterion must be met whenever the instrument is used, all subsequent IV regressions will report the first stage separately as well. The regression includes firm-country and firm-year effects. Standard errors are clustered at the firm-level.

Dependent Variable:	$\Delta \ln(exposure^{long})$		
Model:	(1)	(2)	(3)
Variables			
$\Delta \ln(\widetilde{\tau}^{imp})$	-0.0050***	-0.0049***	-0.0012***
	(0.0003)	(0.0003)	(0.0002)
Fixed-effects			
Firm-Country	Yes	Yes	Yes
Year	Yes	Yes	Yes
Sample	all	extra	intra
Fit statistics			
Observations	200,004	200,004	200,004
Adjusted R^2	0.24479	0.18327	-0.14631

Clustered (Firm) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

To give an example, one possible source of such forces could be changes in investor confidence in the foreign economy by investors in the debt issues, changes in the investor composition, or changes in regulatory regimes coinciding with tariff changes. While these could invalidate the causal interpretation of equation (1.4), as stated before, their effects would have to be present across years for the same firm-country pair, while not being constant. Based on the small relevance of individual firms in trade more broadly (see Figure 1.6), such a violation of the exclusion restriction seems unlikely.

1.4 Results

In this section, the empirical tests outlined in Section 1.3 are conducted using the merged debt-trade-subsidiary sample. First, the cost of borrowing in international markets is compared to that in a firm's home country, followed by a brief discussion of effect heterogeneity and currency choice. Secondly, the firm characteristics influencing the choice to issue internationally are outlined. Finally, a causal link between a firm's physical country exposure and international borrowing are established. Wherever possible, physical exposure will be measured using both the location of subsidiaries, as well as the exposure through trade. As highlighted in Figure 1.4, the underlying sample differs between these analyses, focusing on firms serving a foreign market (exclusively) through subsidiaries and exports, respectively.

1.4.1 Is International Debt Different?

A question that has not yet been answered conclusively is whether debt issued in foreign countries is different from those securities placed in the home market. While Gozzi et al. (2015) are among the first to quantitatively assess the question, their analysis relies on a relatively small sample, limiting the external validity of their findings. It is therefore assessed whether their headline results hold up in the larger sample used by this paper, while also relying on a more robust estimation technique adopted from Benmelech et al. (2022).

Bond Characteristics. A common explanation for why corporations, especially those located in relatively less developed financial markets, choose to issue internationally is that they have outgrown their home market and seek access to deeper financial markets abroad (Allayannis et al., 2003). While prior research has tried to quantify to what extent these incentives play a key role in driving financial markets, it is still unclear whether this is the only, or even predominant, force behind the astonishing growth international debt markets have seen over the last two decades. In this section, domestic and international debt securities, issued by the same entity, are compared based on observables. Recall from on the comparison in means in Table 1.1 that international instruments are on average smaller and more expensive. To alleviate concerns that these differences are driven by sample composition alone, a version of equation (1.1) is going to be estimated, controlling for firm-year effects. Table 1.5 presents the point estimates from regressing various security characteristics on a dummy variable of whether the issuer is a foreign entity. Controlling for a firm-year fixed effect, the analysis is sharpened, by making comparisons between instruments by the same issuer.

Most of the findings in Gozzi et al. (2015) replicate in the significantly larger bond dataset from Refinitiv. International bonds are about 24% larger in size than domestic ones, with tenors that are about 11 months shorter. The latter corresponds to about a quarter of the unconditional average tenor of securities placed in the issuer's home market and is of significant interest, as it hints at the use of the proceeds. If the internal capital markets of firms were frictionless, the funding of any foreign operation should, considering the prior results, be undertaken by issuing debt in the home market and channeling the required capital to wherever it will be deployed. This could be happening for two reasons, either due to a lower cost, or a different benefit associated with borrowing in foreign markets (and thus foreign currencies). Thanks to the more granular data available in this paper, additional variables are evaluated for statistical significance in their difference between domestic and foreign securities. International fixed income instruments are less likely to be senior, while not more or less likely to be senior. Similarly, they are equally likely to be private placements. Interestingly, they are significantly more likely to be callable and less likely to be puttable. These differences indicate that the issuer retains more rights regarding the future adjustment of the amount outstanding for foreign debt securities than they do when issuing in their home market.

Offering Yield. As international debt seemingly differs substantially from securities places in the home market, this naturally begs the question whether firms solely borrow abroad as doing so is cheaper than issuing comparable securities at home. To assess this possibility, equation (1.1) is estimated on the offering yields, i.e., the yield to maturity at issuance, of the securities in the sample. As reasoned previously, the offering yield is a close approximation of the actual cost of borrowing a firm faces. While some may argue that yields in the primary market are biased since only a small number of investors take part in it, this concern is unfounded. While the primary market is indeed 'exclusive,' the yields therein remain the closest measure of the actual cost of

Table 1.5: Other Security Characteristics

This table reports coefficient estimates of regressing different security characteristics on a dummy variable that takes on a value of one if the ultimate parent of the issuing entity is a foreign corporation and zero otherwise. Firm-year fixed effects are included to control for financial conditions. Standard errors clustered at the firm-level are reported in parentheses.

Dependent Variables:	$\mathbb{1}(senior)$	$\mathbb{1}(secured)$	$\mathbb{1}(private)$	$\mathbb{1}(call)$	$\mathbb{1}(put)$	$tenor$	$\ln(face)$
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Variables							
International	-0.1623*	-0.0019	-0.0343	0.0275*	-0.0120***	-10.83*	0.2400*
	(0.0880)	(0.0047)	(0.0738)	(0.0153)	(0.0029)	(6.149)	(0.1441)
Fixed-effects							
Firm-Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics							
Observations	348,296	348,296	348,010	348,296	348,296	348,296	76,311
Adjusted R^2	0.73292	0.78581	0.58907	0.65684	0.77949	0.52501	0.72917

Clustered (Firm) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

debt of the corporate, as it is the time at which the cost of an issuance is determined. To ensure that the comparison is meaningful, the specification is saturated with fixed effects and controls for all major bond characteristics. Estimates are reported in Table 1.6; across all columns the dependent variables are defined in percentage points, i.e., a coefficient of one corresponds to one percentage point, or 100 basis point (bps). Standard errors are clustered at the firm-level.

The dependent variable in columns 1 and 2 of Table 1.6 is the raw offering yield. By including firm-year fixed effects, the comparison between international and domestic issues is made based on groups of debt securities issued by the same firm in the same year but in different jurisdictions. The coefficient on *International* in column 1 corresponds to a difference in yields between international and domestic securities of 35 bps. Considering the unconditional average offering yield of about 6.5% in the sample, this is a sizable difference. It furthermore goes in the opposite direction of the findings is Gozzi et al. (2015), which likely is attributable to much larger representation of bonds issued in developing countries in the sample at hand. Since this effect might be driven by differences between various currencies, a currency-year fixed effect is added to the specification in column 2. While the attenuation of the effect from 35 to 14 bps is consistent with this concern, the difference remains statistically significant and large relative to the sample average.

Having found that international securities differ from domestic ones in their tenor (see Table 1.5), the concern that this wedge in spreads is attributable to differing maturities remains. To address this concern, the dependent variable is replaced with the spread of the offering yield over the corresponding risk-free rate in column 3 and 4. The risk-free rate is obtained from the yield on government securities of the primary issuers in a given currency for the same maturity. Tenors are interpolated following Nelson and Siegel (1987) and Svensson (1994); see Appendix Section A.5 for details. Column 3 mirrors the specification in column 1, including year fixed effects only. The coefficient of interest remains statistically significant at about 14 bps, similar in magnitude to that in column 2. The inclusion of currency-year fixed effects in column 4 does not lead to a further attenuation in the effect. It is evident that the currency-year fixed effect fails to explain any variation beyond the subtraction of r_f. This observation would be consistent with the slope of the yield

Table 1.6: Pricing – Offering Yield

This table reports coefficient estimates for the regression specified in equation (1.1). In columns 1 and 2, the dependent variable is the security's yield to maturity at issuance, whereas in columns 3 and 4 the risk-free rate for the corresponding currency and maturity has been subtracted. The dependent variable is specified in percentage points. Firm-year fixed effects are included in all specifications, to control for time-varying firm effects. Standard errors clustered at the firm are reported in parentheses.

Dependent Variables:	ytm_b		$ytm_b - r_f$	
Model:	(1)	(2)	(3)	(4)
Variables				
International	0.3527**	0.1423**	0.1356**	0.1263**
	(0.1376)	(0.0655)	(0.0596)	(0.0586)
$\mathbb{1}(call)$	0.8817***	1.005***	0.8707***	0.8764***
	(0.2102)	(0.2417)	(0.2538)	(0.2835)
$\mathbb{1}(put)$	-1.694***	-1.572***	-1.648***	-1.524***
	(0.1612)	(0.1391)	(0.1373)	(0.1544)
Tenor	0.0054***	0.0053***	0.0022***	0.0009
	(0.0007)	(0.0009)	(0.0002)	(0.0012)
Fixed-effects				
Firm-Year	Yes	Yes	Yes	Yes
Currency-Year		Yes		Yes
Fit statistics				
Observations	324,422	324,422	291,060	291,060
Adjusted R^2	0.67185	0.73394	0.48742	0.51387

Clustered (Firm) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

curve not varying much over time within a given country.

To assess whether the model is correctly specified, it is worthwhile to also consider the signs of the effects of the remaining control variables. As an embedded call option grants the issuer the right to redeem the security early, it naturally translates into a higher cost of borrowing. The opposite holds for put options. Issues with longer tenors similarly are relatively more expensive, as most yield curves are increasing in maturity. The upshot of the analysis is that international securities command a significant premium over their domestic counterparts. Nevertheless, as the analysis of the primary market relies entirely on actual issuance events it is subject to concerns regarding market timing. This concern will be addressed by using trading yields from the secondary market instead.

Trading Yield. To alleviate concerns that the differences in yields and spreads are driven by the (poor) market timing of issuing firms, the analysis conducted on the offering yield (spread) are repeated on the trading yield (spread) for the part of the sample for which secondary market prices are available. While offering yields correspond to the actual cost of borrowing associated with a debt security that has been issued, secondary market yields have the natural interpretation of the marginal cost of borrowing at the time the market price is observed. Naturally, the sample of instruments for which secondary market yields are available is a selected one, with issues being larger and having lower yields than the average security (see Table 1.1). Nevertheless, the analysis can address the main concern that results on yields and spreads in Table 1.6 are driven by market timing rather than capturing the contemporaneous difference in borrowing costs between domestic and international issues.

The estimates for equation (1.1) with trading yields as dependent variable are reported in Table 1.7. Trading yields are collected from different trading venues to construct a meaningful time-series. To ensure robustness to liquidity effects, which would lead to wider bid-ask spreads, the mid yield is used in all regressions. As before, the dependent variable in columns 1 and 2 is the yield, from which the contemporaneous risk-free rate r_f in the same currency and for the same maturity is subtracted in columns 3 and 4. From the prior analysis it is clear that controlling for currency-year fixed effects in column 2 is very important, as yield curves differ significantly across countries. This gives rise to the preferred specification, delivering an estimate of the difference in yields of 11 bps, slightly lower than the difference in yields in the primary market. Moving to spreads, the general magnitude of the results once again resembles those in Table 1.6, at 19 bps in column 3 without currency-year fixed effects, which are attenuated to 12 bps in column 4.

A quick look at the remaining control variables reveals that the coefficients are predominantly insignificant. The exceptions are the dummy turning on for securities with embedded put options, which demand significantly lower yields, as one would expect. The coefficient on *Tenor* is similar

Table 1.7: Pricing – Trading Yield

This table reports coefficient estimates for the regression specified in equation (1.1). In columns 1 and 2, the dependent variable is the security's yield to maturity in the secondary market, whereas in columns 3 and 4 the risk-free rate for the corresponding currency and maturity has been subtracted. The dependent variable is specified in percentage points. Firm-year fixed effects are included in all specifications, to control for time-varying firm effects. Standard errors clustered at the firm are reported in parentheses.

Dependent Variables:	ytm_b		$ytm_b - r_f$	
Model:	(1)	(2)	(3)	(4)
Variables				
International	-0.0052	0.1104***	0.1970***	0.1206***
	(0.0921)	(0.0273)	(0.0364)	(0.0283)
$\mathbb{1}(call)$	0.0623	-0.0053	0.0754	0.0037
	(0.1288)	(0.1071)	(0.1112)	(0.1025)
$\mathbb{1}(put)$	-0.8912**	-0.8649**	-1.137***	-0.9455***
	(0.3860)	(0.3566)	(0.4134)	(0.3504)
$\mathbb{1}(senior)$	-0.0034	0.0387	0.1795	0.0440
	(0.1589)	(0.1454)	(0.1461)	(0.1685)
$\mathbb{1}(secured)$	-0.0448	-0.1402	-0.1539	-0.1264
	(0.2020)	(0.1940)	(0.1902)	(0.1936)
Tenor	0.0043***	0.0039***	0.0018***	0.0018***
	(0.0003)	(0.0003)	(0.0002)	(0.0002)
Fixed-effects				
Firm-Year	Yes	Yes	Yes	Yes
Currency-Year		Yes		Yes
Fit statistics				
Observations	1,041,567	1,041,567	1,034,835	1,034,835
Adjusted R^2	0.61900	0.63536	0.58323	0.58574

Clustered (Firm) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

in magnitude and significance as in the prior set of results, consistent with upward-sloping yield curves. Overall, the model seems to be properly specified, revealing a meaningful difference in spreads between domestic and foreign securities, which cannot be reconciled by the most salient characteristics.

The remarkable similarity in coefficients across Table 1.6 and Table 1.7 indicates that the results are not driven by differences in market timing between domestic and foreign debt issuance. This result differs from that in Gozzi et al. (2015), who find that international issuances are cheaper than domestic ones. This difference in conclusions warrants a brief note on the relevance of default risk across securities. In the regressions in this paper, default risk is not explicitly controlled for, potentially explaining part, if not all, of the difference in spreads. This is a deliberate choice, as doing so would be misguided given the question at hand. As the interest in this paper lies on the firm's rationale for issuing abroad, any difference in borrowing costs attributable to differential default risk would still fall back onto the borrower in the form of higher spreads. The difference in spreads also tells a different story than that in Liao (2020). While firms may exploit deviations from CPI between the USD and the EUR, the results in this paper are predominantly driven by (positive) yield differences between developed and developing markets, rather than between two developed markets.

Heterogeneity. The estimates on the differences in spreads inherently are a (variance-weighted) average. To assess the heterogeneity of the effect, the analysis on the primary market yields is repeated in groups defined by (a) the country of issuance, (b) the firms' location, (c) location-country pairs, and (d) within the firms themselves. The coefficients from these regressions have the following, natural, interpretations: (a) differences in spreads between issuances of domestic and foreign firms in the same market, (b) differences in spreads between domestic and foreign issuances by firms located in the same country, (c) difference in spreads between domestic and foreign issuances in country c for firms located in country d, and (d) difference in spreads between domestic and foreign issuances for a given firm. Each of these tests captures a distinct notion of heterogeneity, the distribution of the coefficients are presented in Figure 1.7. As can be gathered from the distributions, the mode of all distributions is very close to zero or slightly positive. All distributions are on average positive, consistent with the previous regressions. As in Gozzi et al. (2015), the effect is primarily driven by issuance in developing countries by firms themselves located in developed economies.

Currency Choice. Last but not least, as the existing literature places a lot of emphasis on curren-

Figure 1.7: Pricing – Heterogeneity

The figure below depicts the heterogeneity in the differences in spreads between domestic and foreign fixed income securities. Four different analyses are conducted to get at different notions of heterogeneity. The sample is split into groups based on (a) the country of issuance, (b) the firms' location, (c) location-country pairs, and (d) the firms themselves. The coefficients from these regressions have the following interpretations: (a) differences in spreads between issuances of domestic and foreign firms in the same market, (b) differences in spreads between domestic and foreign issuances by firms located in the same country, (c) difference in spreads between domestic and foreign issuances in country c for firms located in country d, and (d) difference in spreads between domestic and foreign issuances for a given firm.

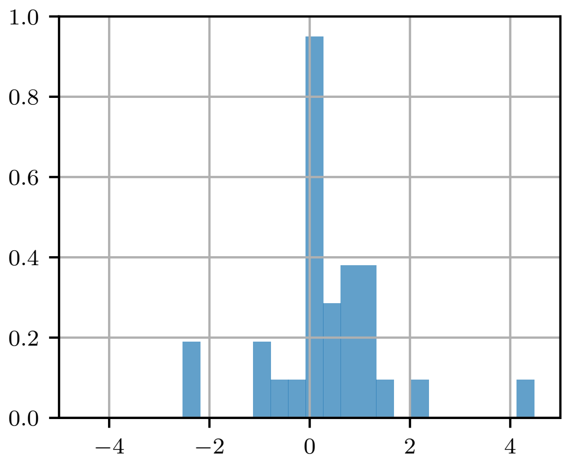

(a) Country of Issuance

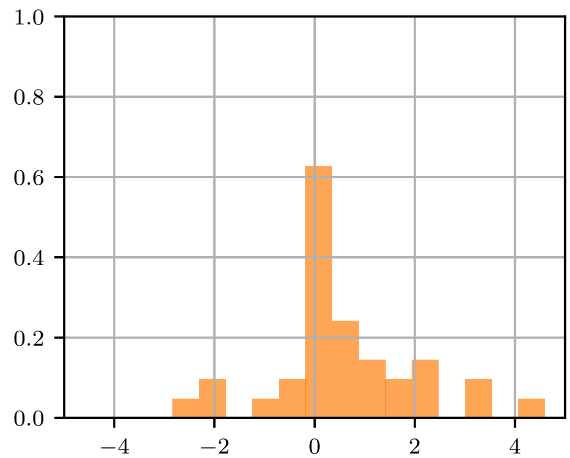

(b) Country of HQ

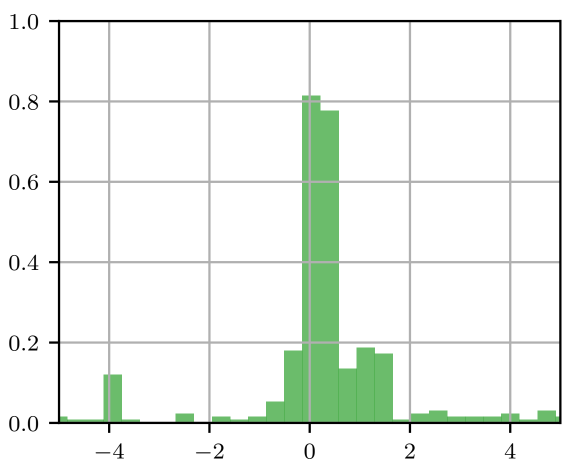

(c) Country Pairs

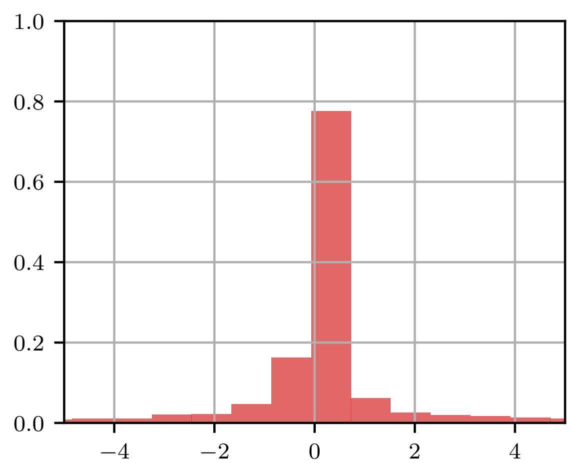

(d) Firm

Figure 1.8: Currency Choice

The figure below depicts the share of total issuance in country c by domestic (international) firms that are in the country's primary currency. Both domestic (orange circles) and international (blue squares) entities issue predominantly in local currency units (LCU).

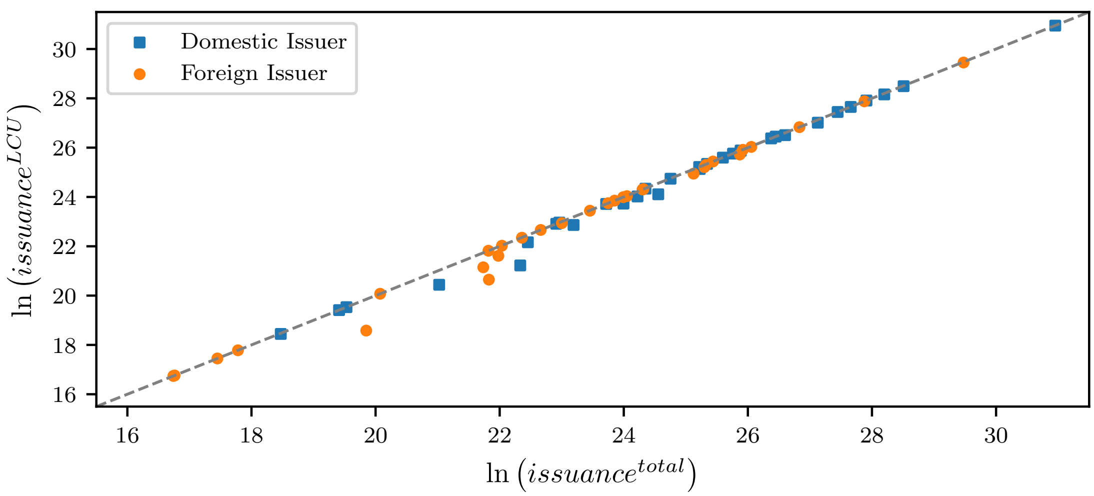

cies (Keloharju and Niskanen, 2001), the differences in currency denomination of the securities in the sample warrants a brief discussion. As seen in Table 1.1, the majority of all issues are in the currency of the market in which they are placed. While this indicates that, to a first order, the notion of country in this sample corresponds to the currency of issuance, it is important to briefly check whether there are any significant differences between countries in this observation. The fact that this is not just driven by securities placed in the U.S. can be seen from Figure 1.8, which depicts the share of securities denominated in the local currency for all markets in the sample.

While foreign issuers are more likely than domestic issuers to issue debt in a currency other than the local currency of the market of issuance, both issue predominantly in the local currency. This is not surprising for several reasons. Firstly, untabulated results suggest a significant penalty when issuing in foreign currencies. Therefore, firms are likely better off issuing in local currency and swapping the associated payment stream into their local currency. Secondly, issuance in foreign countries, as argued in greater detail below, is one of the major ways of obtaining funding in foreign currencies. Especially in light of the additional spread many corporations face by tapping foreign capital markets, doing so likely is the cheapest way of obtaining a liability in the desired currency.

1.4.2 Who Issues Internationally?

Having shown that there is a significant premium associated with raising funds in foreign capital markets for a lot of companies, this evidently poses the question who these entities are. In contrast to prior research (Gozzi et al., 2015; Keloharju and Niskanen, 2001), the data in this paper allows for a within-firm, cross-country analysis of the driving forces behind issuing internationally. The evidence so far is consistent with firms using international capital markets to raise capital in foreign currencies. Naturally, one would expect the choice of issuance to be closely related to invoicing currencies (Boz et al., 2022; Gopinath et al., 2010). Unfortunately, data at the firm-level on invoicing currencies is not readily available. However, data on the location of subsidiaries and the composition of sales can serve as proxies for the composition of a firm's currency exposure. Relating back to the different modes of market entry in Figure 1.4, there are two sources of exposure originating from international operations; (1) stemming from long-term projects, e.g., the establishment of a foreign subsidiary, or (2) related to short-term capital needs arising from exports (Allayannis et al., 2003). If firms jointly choose their operational and financial exposure to foreign countries, it can be hypothesized that firms with greater relative exposure to a foreign market have a stronger incentive to issue there. In this section, the hypothesis that firms with a significant *physical exposure* to a given country are the ones tapping the corresponding capital market is put to the test.

Subsidiaries. The first set of analyses is focused on firms that have international operations. Table 1.8 presents coefficient estimates from regressing the logarithm of the value of debt outstanding in a given country on the logarithm of the number of subsidiaries a firm has in the same country. In column 1, the total number of subsidiaries is considered. Columns 2 and 3 restrict the count to manufacturing and financial subsidiaries, respectively. All specifications include three granular layers of fixed effects. Firm-country effects ensure that the identification comes from variation in the number of a firm's subsidiaries in a given country. Firm-year fixed effects control for a firm's financial condition and other time-varying unobservables. Lastly, country-year fixed effects absorb

Table 1.8: Propensity to Issue – Subsidiaries

This table presents the coefficient estimates from regressing the amount of debt outstanding in a given country on the natural logarithm of the number of establishments the issuer maintains in the same jurisdiction. The independent variables are broken up into the total number of subsidiaries, the number of manufacturing establishments, and the number of financing subsidiaries. All regressions include firm-country, firm-year, and country-year fixed effects. Standard errors are clustered at the firm-level.

Dependent Variables:	$\ln(exposure^{fin})$		
Model:	(1)	(2)	(3)
Variables			
$\ln(subs_{tot})$	0.2581***		
	(0.0253)		
$\ln(subs_{cor})$		0.2550***	
		(0.0249)	
$\ln(subs_{fin})$			0.6774**
			(0.0531)
Fixed-effects			
Firm-Country	Yes	Yes	Yes
Firm-Year	Yes	Yes	Yes
Country-Year	Yes	Yes	Yes
Fit statistics			
Observations	1,617,677	1,617,677	1,617,677
R^2	0.8970	0.8970	0.8971

Clustered (Firm) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

any time-varying effect within a country, such as policies attracting foreign direct investment.

The coefficients in all three specifications are positive and significant. The baseline delivers an estimated elasticity of 0.2581. A one standard deviation increase in the number of subsidiaries is associated with a 21% increase in the amount of debt issued in the corresponding country. These regressions exclude the home country for all firms in the sample. A major concern with looking at the subsidiary counts and how these numbers relate to the amount of debt issued lies in the importance of financing subsidiaries in issuing securities. To tap a foreign debt market, firms oftentimes are required to set up a local establishment. To alleviate the concern that such effects are driving the results, all subsidiaries that are identified as having been established for financing purposes are removed from the sample in column 2. The coefficient remains basically unchanged, addressing such concerns to some degree. Lastly, to highlight the endogeneity issue between the number of financing subsidiaries and debt issuance, the independent variable in column 3 counts only such subsidiaries. The much greater coefficient is consistent with such concerns being relevant. In the last section of this paper, the opening and closing of subsidiaries in the time-series will be related to changes in the amount of debt outstanding to further tighten the identification.

Exports. A novel measure of a firm's exposure to a country proposed is this paper is the value of (net) exports. Deriving this measure from detailed customs data is one of the key contributions of this paper. In a first pass at the question whether firms jointly choose where to raise capital and which countries to export to, the propensity to issue in country c as a function of the trade surplus/deficit is analyzed. This is accomplished by regressing a dummy variable capturing whether a firm had at least one bond outstanding in a given country on the net exports. A binned scatter plot of the corresponding regression coefficients is presented in Figure 1.9. The regression is run separately for net importers, i.e., net long firms, and net exporters, i.e., net short firms, as the net exposure is not well defined around a value of zero. The coefficient estimates from a regression of issuance on the long/short exposures is presented in Table 1.9.

The data underlying the plot in Figure 1.9 is constructed exclusively from exports to (imports from) unaffiliated entities. Firms with subsidiaries in the focal country are excluded to ensure that only pure exporters are considered. Furthermore, the regressions include controls for firm size as measured by the logarithm of total assets. As large firms are more likely to issue in foreign capital markets, this size effect would otherwise drown out any heterogeneity arising due to differential trade exposures. As the regressions for firms that are net long/short are estimated separately, the two sides of the plot can be interpreted independently. The first observation that can be made concerns the level of the dependent variable between the two sides; firms with a net long exposure are more likely to issue in a given country than those with a net short exposure. Furthermore, while net importers seem not to issue, regardless of the size of their short exposure, the propensity to issue increases with greater long exposure for net exporters.

Figure 1.9: Propensity to Issue – Total Trade

This figure depicts the propensity of firm i to issue debt securities in country c in year t based on the logarithm of their net exposure. Net exposure is computed as $\ln(net\ exposure)_{i,c,t} = \text{sgn}(imp_{i,c,t} - exp_{i,c,t}) \times \ln\left(1 + |imp_{i,c,t} - exp_{i,c,t}|\right)$. Only shipments to/from unaffiliated entities are considered. Blue dots correspond to net importers, i.e., firms that have a net long exposure, while orange ones are net exporters, i.e., firms with a net short exposure. Controls include the natural logarithm of a firm's total assets to control for heterogeneity in firm size.

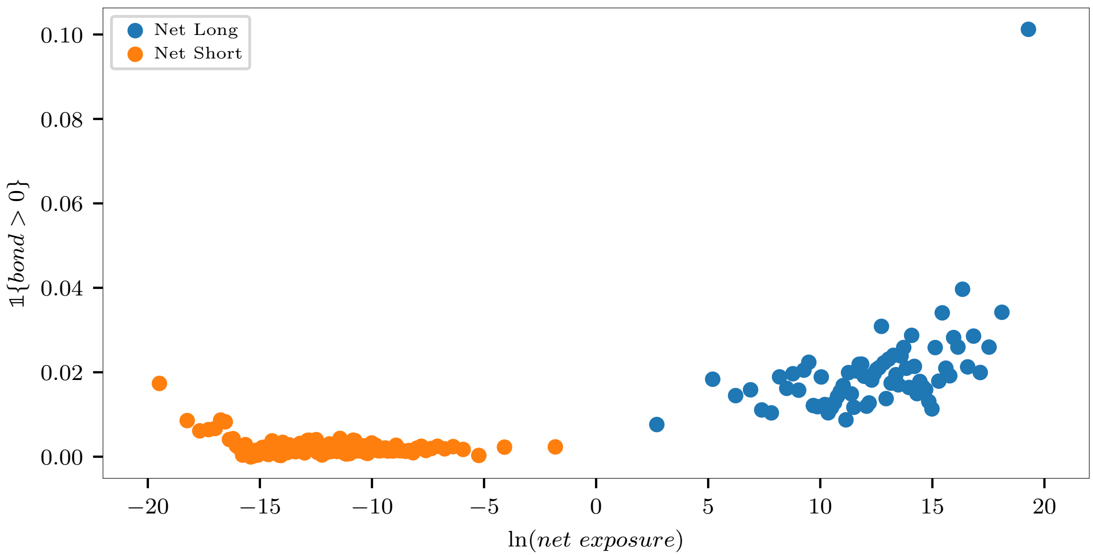

To further control for other confounding factors, the same analysis is cast in a regression framework in Table 1.9. The regression specification is presented in (1.2). The dependent variables in columns 1 through 3 is an indicator variable for whether a firm f has debt outstanding in country c at time t. In columns 4 through 6 this is replaced with the logarithm of the value of financial instruments outstanding in country c. All specifications include firm-year fixed effects and standard errors clustered at the firm-level. The explanatory variable in columns 1 and 4 is a dummy for whether the firm had a net long exposure. Consistent with the observations from Figure 1.9, firms with net long exposures are significantly more likely to issue (see column 1) and correspondingly have larger dollar amounts outstanding (column 4). The coefficient of 0.0224 in column 1, while small, is similar in magnitude to the unconditional propensity to issue debt in foreign jurisdictions.

In columns 2 and 5 the sample is restricted to firms with long exposures and the explanatory variable is replaced with the logarithm of the net exposure. The coefficients in either case are

Table 1.9: Propensity to Issue – Trade

This table reports the coefficient estimates from regressing variables that capture the financial exposure of firm *i* to country *c* in year *t* on measures of the operational exposure. In columns 1 through 3 the dependent variable is a dummy which takes on a value of one if the firms have securities outstanding in country *c* and zero otherwise. The logarithm of the financial exposure is the dependent variable in the remaining columns. The explanatory variable in columns 1 and 4 is an indicator for a net long exposure. Columns 2 and 5 restrict the sample to firms with a net long exposure, whereas the sample in columns 3 and 6 is limited to firms that are net short. Across all specifications the sample is restricted to firms without manufacturing subsidiaries in country *c*. The independent variable in these columns is the logarithm of the net exposure. All regressions include firm-year fixed effects. Standard errors are reported in parentheses and clustered at the firm-level.

Dependent Variables:	$\mathbb{1}(bond)$			$\ln(exposure^{fin})$		
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
$\mathbb{1}(net\ long)$	0.0224***			0.4176***		
	(0.0030)			(0.0570)		
$\ln(net\ exposure)$		0.0043***	-0.0004***		0.0882***	-0.0069***
		(0.0011)	(0.0001)		(0.0232)	(0.0023)
Fixed-effects						
Firm-Year	Yes	Yes	Yes	Yes	Yes	Yes
Sample	all	long	short	all	long	short
Fit statistics						
Observations	232,819	47,476	185,343	232,819	47,476	185,343
R^2	0.46252	0.77349	0.51876	0.46137	0.76487	0.51651

Clustered (Firm) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

positive and highly significant. These findings are consistent with the increase in propensity to issue the larger the net exposure. Furthermore, they highlight that this effect is not just present at the extensive but also the intensive margin. Moving to columns 3 and 6, which restrict the sample to firms with net short exposures. The coefficient estimates are negative. Since a short net exposure has a negative sign, they are to be interpreted similarly to those in the previous column, i.e., the larger the short exposure the stronger the effect. The coefficients for firms that are net short are much smaller in absolute value than those for the net long ones, consistent with the significantly less pronounced effect on the short side in Figure 1.9.

The findings in this section suggest two things. On the one hand, firms seemingly finance foreign subsidiaries with debt raised in the country in which they will operate. This is consistent with an asset-liability-matching motive. Specifically, as they are likely to face operating costs and associated revenues in the foreign currency, using debt denominated in the same shields them from any FX exposure beyond that of translating profits. Similarly, for exporters the findings are highly consistent with firms using foreign debt markets to hedge FX exposures arising due to their international operations. In the following section the merit of these hypotheses is going to be assessed empirically.

1.4.3 Currency Hedging

This section will propose and provide empirical evidence in support of the hypothesis that firms use international debt markets as a means of hedging the FX exposure associated with their international activities. To make the case for this hypothesis it is important to briefly subsume the findings up to this point. International debt securities seemingly are an expensive form of borrowing when compared to those issued by the same firm in their home market. The fact that, in equilibrium, firms still issue internationally seems to suggest that there are tangible benefits associated with doing so. Moving on to issuance choice, firms that issue internationally have one of two possible motives. On the one hand, among those firms with international subsidiaries, debt is raised in the country, and thus the currency where subsidiaries are located. On the other hand, pure exporters are disproportionately more likely to raise capital in those countries they are exporting to. This effect is (almost) entirely driven by firms which export more goods to the foreign country than they import from it, leaving them with a net long exposure. Below, the following will be discussed in turn; (1) a firm's hedging motive that can rationalize these findings, (2) an analysis on firms with subsidiaries lending support to the hypothesis, and (3) an analysis on exporters, which delivers a causal estimate of the strength of the mechanism.

Mechanism. As argued in the previous sections, it is natural to hypothesize that multinational firms use international financial markets to hedge their exposure to foreign currencies. There are two natural ways in which exposure to foreign currencies can arise. Firms that have international operations, i.e., subsidiaries in foreign jurisdictions, usually incur expenses in foreign currency units (FCU), e.g., wages and other costs associated with manufacturing in a different country. Furthermore, at the time a subsidiary is set up, an upfront investment will be required. It is beneficial to fund foreign subsidiaries with local debt and, instead, only worry about the conversion of future surplus into local currency units (LCU). This method of dealing with currency and country risk is referred to as *asset-liability matching*. There is ample anecdotal evidence that firms actively engage in this behavior to manage their FX exposure. Furthermore, by raising funds in foreign countries, international operations can be insulated from the parent company. While selective default is not observable in the fixed income data, having the option to abandon a foreign subsidiary by handing it over to (local) creditors may be of substantial value to multinational corporations.

Firms, which are pure exporters, face similar, albeit more nuanced, considerations. Moving goods internationally ties up a large amount of capital, which can be financed in a variety of ways (Antràs and Foley, 2015). Alternatively, borrowing in international debt markets can achieve a similar goal. Figure 1.10 depicts how the creation of such liabilities can be used to hedge any exchange rate exposure. A firm that sells goods to foreign customers, by assumption invoicing them in FCU, generates an FCU receivable. Due to long international shipping times, the receivable will only lead to a cash inflow once the goods have arrived at the customer. In the meantime, exchange rates move giving rise to FX risk. To hedge this exposure, the firm can borrow in the foreign market, again denominated in FCU, and convert the proceeds to LCU at the spot rate. Having borrowed in the fixed income market, the receivable can be used to pay coupons and the principal. This way of alleviating FX risk is called a *natural hedge*. It is similar to the notion of asset-liability matching above.

A natural question then becomes why corporations would want to engage in either asset-liability matching or using natural hedged when financial instruments aimed at managing such exposures exist, e.g., options and futures. Evidence suggests that for the large corporations in

Figure 1.10: Natural Hedge

This figure depicts the cash flows involved in a natural hedge. The exporter expects cash inflows in foreign currency units (FCU). To hedge the exposure, a liability is created against, by borrowing in FCU. The proceeds are converted at the spot rate to neutralize the exposure, which the future FCU inflows are used to pay off the interest and principal. Adapted from Eiteman et al. (2015).

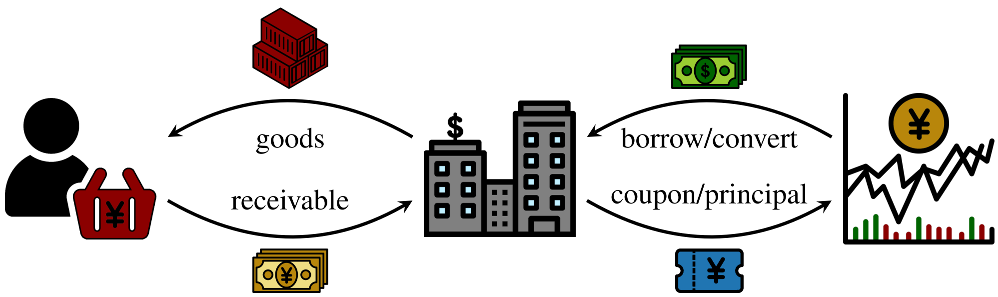

the sample, when trying to hedge a significant currency exposure, the cost of using derivatives is prohibitive. While this should be fairly intuitive in the case of (large) foreign subsidiaries, many exporters have receivables that would require entering and managing derivative positions with many millions of notional. While options require an upfront payment, maintaining a large futures position requires a potentially large commitment of collateral. Tapping capital markets, on the other hand, is often a significantly cheaper way of neutralizing FX exposure. Issuing through a known underwriter oftentimes is much easier. Having spoken to many practitioners working in debt capital markets, a common narrative is that underwriting a new issuance for an existing client boils down to a mere change of numbers in existing documents followed by a few phone calls to parties that are known to be interest in such an investment. It therefore seems likely that firms indeed use international debt markets in this way.

Both *asset-liability matching* (for multinationals) and *natural hedging* (for exporters) have testable implications. In fact, the underlying mechanism is almost identical, with the former being related to the balance sheet, while the latter is unique to the income statement. Both of them are thus assessed empirically leveraging the framework outlined in Section 1.3.

Asset-Liability Matching. To test whether firms adjust their borrowing following the opening and closing of subsidiaries, equation (1.4) is estimated on the data of multinational corporations. The coefficient estimates are reported in Table 1.10. While the endogeneity concerns this regression is subject to are less severe when applying it to the number of subsidiaries as measure of operating exposure, they do remain. Due to the absence of an instrument for the opening and closing of establishments, the coefficient estimates should be interpreted with caution and not be considered causal effects. The estimates suggest an elasticity of issuance with respect to the firm's local footprint of between 0.11 and 0.12 (columns 1 and 2). For the reasons outlined before, it is important to exclude financing subsidiaries, as they are mechanically related to the issuance of securities. The preferred specification in column 2 controls for three sets of fixed effects. Firm-country fixed effects absorb any time-invariant effects driving a firm's choice to operate in a given country. Firm-year effects control for the firm's time-varying financial condition. Country-year effects account for changes in economic conditions in the country, which are present across firms. Standard errors are clustered at the firm-level. The effect size of 0.1084 suggests that a 1% increase in the number of subsidiaries is associated with an increase in the face value of debt outstanding of 10.84%.

The natural question to ask is why the effect is smaller than 1. While one would expect the relationship between borrowing and the number of subsidiaries to be one-to-one if all subsidiaries were identical, this clearly is counterfactual. Specifically, since most of the entities in the sample adjust their operating footprint on the intensive margin, i.e., open (close) marginal subsidiaries without entering (leaving) the country, it is expected that the changes in subsidiary count are driven by relatively smaller establishments. Unfortunately, the data does not offer sufficient granularity to control for the different sizes of subsidiaries. Nevertheless, the size of the coefficient is in line with this interpretation. Another important, related, consideration are the differences in financing needs for marginal establishments. Even if all establishments were the same size, it likely is the case that initial market entry requires significantly more resources than the opening of a (marginal) subsidiary. These teething issues would be another natural explanation for the magnitude of the coefficient. Despite these concerns, the positive association between establishment counts and debt outstanding are in line with the ex-ante expectations.

Table 1.10: Intensive Margin – Subsidiaries

This table presents coefficient estimated from regressing the log-change in nominal value of securities outstanding on the change in the number of a firm's subsidiaries in the same country. Standard errors clustered at the firm-level are reported in the parentheses.

Dependent Variables:	$\Delta\ln(debt)$		
Model:	(1)	(2)	(3)
Variables			
$\Delta\ln(subs_{tot})$	0.1213***		
	(0.0126)		
$\Delta\ln(subs_{cor})$		0.1084***	
		(0.0124)	
$\Delta\ln(subs_{fin})$			0.3047**
			(0.0268)
Fixed-effects			
Firm-Country	Yes	Yes	Yes
Firm-Year	Yes	Yes	Yes
Country-Year	Yes	Yes	Yes
Fit statistics			
Observations	1,410,946	1,410,946	1,410,946
R^2	0.3724	0.3724	0.3725

Clustered (Firm) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

It cannot be ruled out that there is endogeneity between the establishments of subsidiaries and debt issuance. For instance, even absent the desire to fund foreign operations with capital obtained from foreign markets, local market conditions could lead to a firm issuing securities in a given country, while also opening factories due to an increase in demand in the foreign economy. While the country-year fixed effect controls for such effects that are constant across firms, heterogeneous effects across industries could still be at play. Since controlling for such effects is not possible in the establishment data, these results should be considered primarily expositional. In the following section the trade data is used to estimate a causal relationship between a firm's exposure and decision to borrow in a foreign capital market.

Natural Hedge. If firms use international capital markets to hedge the FX exposure originating from their operations, there should be a significant relationship between changes in a firm's country

Table 1.11: Intensive Margin – All Trade

This table presents coefficient estimated from regressing the log-change in nominal value of fixed income instruments outstanding on the change in imports. In columns 2 and 3 the change in imports is instrumented using the tariff instrument. Standard errors clustered at the firm-level are reported in the parentheses.

	OLS	2SLS	
Dependent Variables:	$\Delta\ln(bond)$	$\Delta\ln(imp)$	$\Delta\ln(bond)$
Model:	(1)	(2)	(3)
Variables			
$\Delta\ln(imp)$	0.0004		0.0497***
	(0.0006)		(0.0188)
$\Delta\ln(\widehat{\tau})$		-0.0017***	
		(0.0002)	
Fixed-effects			
Firm-Country	Yes	Yes	Yes
Year	Yes	Yes	Yes
Fit statistics			
Observations	317,811	233,953	233,953
F-test (1st stage)		74.70	

Clustered (Firm) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

exposure and the borrowing in that jurisdiction. To assess the validity of this channel, (1.4) is estimated on the sample of exporters, employing the instrumental variable strategy outlined in Section 1.3. The OLS and 2SLS estimates are presented in Table 1.11. Since long exposures originating from imports are the driving force behind issuance (see Figure 1.9 and Table 1.9), imports are the main focus of this analysis. In column 1, the effects from equation (1.4) without using the instrument are reported. Due to the low propensity to issue abroad, the estimate is inherently biased towards zero. Furthermore, there remains a lot of noise in the shipments, another concern that the instrumental variable (IV) strategy addresses to some extent.

Columns 2 and 3 present the parameter estimates from employing the IV strategy discussed in Section 1.3. As in column 1, the specifications include firm-country as well as year fixed effects, while the standard errors are clustered at the firm-level. The first stage delivers a significant coefficient of -0.0017, which is comparable to results in the literature when adjusting for the slight

Table 1.12: Intensive Margin – Extra-Firm Trade

This table presents coefficient estimated from regressing the log-change in nominal value of securities outstanding on the change in imports. In columns 2 and 3 the change in imports is instrumented using the tariff instrument. The trade exposure variable is constructed from flows between unaffiliated entities only. Standard errors clustered at the firm-level are reported in the parentheses.

	OLS	2SLS	
Dependent Variables:	$\Delta\ln(bond)$	$\Delta\ln(imp_{extra})$	$\Delta\ln(bond)$
Model:	(1)	(2)	(3)
Variables			
$\Delta\ln(imp_{extra})$	0.0007		0.0508***
	(0.0006)		(0.0194)
$\Delta\ln(\widehat{\tau})$		-0.0017***	
		(0.0002)	
Fixed-effects			
Firm-Country	Yes	Yes	Yes
Year	Yes	Yes	Yes
Fit statistics			
Observations	317,811	233,953	233,953
F-test (1st stage)		60.25	

Clustered (Firm) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

differences in how the measure of changes in tariffs is constructed. The first stage F-statistic has a value of 74.70, well in excess of the critical value of 10 (Stock et al., 2002). A negative coefficient on the instrument is in line with economic theory, as higher tariffs hinder international trade. Moving on to column 2, in which the change in financial exposure is regressed on the instrumented change in trade exposure, the coefficient is positive and significant. This is consistent with the ex-ante expectations and in line with the results in Table 1.10. The magnitude of 0.05 translates a one standard deviation shock to the physical exposure to a change in the amount of debt outstanding by about 10 percent. The effect is not just statistically significant, but also economically meaningful. Since a not insignificant share of the sample does not issue debt in the affected countries, this local average treatment effect (LATE), while still sizable, is biased towards zero and thus a lower bound on the causal effect.

While the IV strategy overcomes issues of endogeneity between trade and financial exposure, the measure of trade exposure used in Table 1.11 pools shipments to affiliated as well as unaffiliated companies. Since the customs values of the former might not be meaningful due to strategic transfer pricing, the analysis is repeated in Table 1.12, removing any such shipments in the construction of the endogenous variable. The results remain almost entirely unaffected, suggesting that the tariffs instrument works for both intra-company and extra-company shipments. The OLS estimate in column 1 is larger but equally insignificant as in Table 1.11. When moving to the 2SLS estimates in columns 2 and 3, the first stage remains entirely unchanged, while the coefficient in the second stage is insignificantly larger.

Overall, the results are consistent with firms using international capital markets in managing their exposure to foreign currencies which arises from their operations. This link can be established for different measures of operating exposure, consistent with two highly related mechanisms, *asset-liability matching* and *natural hedging*.

1.5 Conclusion

International debt markets have grown in importance over the past two decades. This paper sheds light on the securities issued therein, the entities issuing them, and their motivations for doing so. The existing literature has focused on the decision of issuing internationally and/or in foreign currencies as a binary one, abstracting away from the choice of market/currency. This paper fills this gap, by leveraging a novel dataset combining the universe of fixed income securities, detailed information on the location as well as nature of firm subsidiaries, and the complete customs records for 17 developing and developed economies. The sample period has been chosen as 2007-2022 to jointly maximize the breadth and depth of the data.

The analyses in this paper seek to address three crucial questions; (1) how does debt issued abroad differ from instruments issued domestically, (2) which firms choose to issue internationally, and (3) what links a firm's international exposure on the funding and operating side. With respect to the first, detailed comparisons between similar securities issued by the same entities in the same year reveals that international fixed income instruments carry higher yields, have larger nominal amounts, and shorter tenors. The increase in spreads of about 12 bps when issuing abroad is inconsistent with firms solely doing so to minimize borrowing costs and suggests that there are other tangible benefits attached to borrowing internationally. Turning to the entities that issue abroad, two key determinants of international issuance are found. Firstly, firms with international operations issue in exactly those foreign countries where they have manufacturing subsidiaries. Secondly, among the countries that are served through exports, debt is raised in those markets to which the firm exports more than it imports from. The larger the difference between the two, the more likely a corporation is to issue.

These two findings are consistent with firms using international debt markets to engage in *asset-liability matching* when committing to long-term projects abroad and engaging in *natural hedging* in an attempt to manage the exchange rate exposure originating from exporting. These mechanisms are not mutually exclusive. In an empirical assessment of the mechanism, a significant relationship between changes in the operating and financial exposure is uncovered. In case of exporters, the relationship is a causal one, as it is identified using an instrumental variable strategy. More concretely, following the trade literature, changes in tariffs are used to instrument for changes in a firm's exposure to foreign countries through trade.

The findings in this paper are consistent with anecdotal and prior empirical evidence on the firm's choice where to raise capital. They add to the literature by causally identifying this link and quantifying the relationship between a firm's operating and financing exposure. In light of the seemingly unstoppable growth in international trade, they provide a micro-foundation for the contemporaneous growth in cross-border borrowing. Future research should place further emphasis on the intersection of international finance and trade from the firm's perspective, an area that still lacks additional empirical evidence on the driving forces.

Chapter 2:
Did You Catch the Game Last Night?
Peer Group Effects in Sell-Side Analyst Forecasts[1]

2.1 Introduction

An ever growing body of research suggests that interactions between peers have meaningful implications for financial decision making. In early work, Madrian and Shea (2001) and Duflo and Saez (2003) show that an individual's decision to participate in certain employer-sponsored retirement plans is affected by the choices of their co-workers. Similarly, Hong et al. (2004) find that investors find the stock market more attractive when a larger number of their peers participate. Bailey et al. (2018a,b, 2022) proxy for peer groups through data from Facebook and tease out the relevance of peer effects in the housing market and product adoption. The relevance of selective exposure to peers (echo chambers) even among sophisticated investors is the subject of a recent study by Cookson et al. (2023).

Despite the overwhelming evidence regarding the relevance of peer effects, ambiguity remains around the exact mechanisms at play. On the one hand, individuals draw on their social networks to obtain additional, meaningful information to base their beliefs on (Bailey et al., 2022; Cookson et al., 2023; Fischer, 2022). On the other hand, peers may influence the sentiment and emotional aspect of an individual's worldview, which can in turn shape their behavior.[2] In a similar vein, the issue of homophily is a pervasive concern in the estimation of peer effects (Angrist, 2014).

Differentiating between these mechanisms is difficult, as it is often unclear whether individuals are able to parse out spurious information when drawing on their social network, and it is rare to

[1]This chapter is based on Fischer and Shore (2023). We thank seminar participants at Columbia University.

[2]There exists a vast literature in psychology on peer influence (in adolescents). Brechwald and Prinstein (2011) provides an excellent summary.

see an instrument employed when assessing peer effects. While Bailey et al. (2018a,b) suggest that households react to the experiences of their peers, it remains an open question whether doing so is rational. In their setting, the US housing market, one can argue that the experiences of peers is a meaningful source of information as long as there are common factors across regions. Furthermore, while households are generally considered unsophisticated investors, the question of how professionals draw on information from their peers remains an open one.

Several papers have considered how professionals react to information. Similar to our setting, Kempf and Tsoutsoura (2021) consider analysts as financial professionals, and attempt to identify how they incorporate outside information into their decision making. Specifically, they analyze the reaction of credit analysts to changes in the 'color' of the white house depending on partisanship. However, once again it is difficult to argue that these responses are spurious, as Democrats will likely expect the policies of a Democratic president to be economically more beneficial.

Drawing on the existing literature on the non-trivial impact of sports results on decision making (Edmans et al., 2007; Eren and Mocan, 2018), we exploit a novel source of plausibly exogenous variation in equity analyst forecasts that is well-suited to analyzing the sentiment component of local diffusion: the shock to an analyst's sentiment resulting from their college team winning/losing the NCAA National Championship Football game.

We choose this setting for two reasons. Firstly, equity analysts operate in a high stakes setting, with numerous papers demonstrating the significant implications that their forecasts can have on firm and stock market behavior.[3] Yet evidence is widespread that their forecasts are subject to influence from plausibly arbitrary sources.[4] As such, it is not far fetched to claim that college foot-

[3]Beginning with early work on the informational content of analyst forecasts (Elton and Gruber, 1972; Fried and Givoly, 1982), several papers have documented profitable trading strategies that are based on forecasts and their revisions (Elton et al., 1981; Givoly and Lakonishok, 1979, 1980; Griffin, 1976; Imhoff and Lobo, 1984). More recent work has examined stock market reactions to forecast revisions (Frankel et al., 2006; Gleason and Lee, 2003), the overweighting of forecasts by investors (So, 2013), the impact of analyst coverage on crash risk (Kim et al., 2018), and firms' earnings management responses to forecasts (Almeida et al., 2016; Bhojraj et al., 2009; Shore, 2023; Terry, 2015).

[4]Factors that have been shown to influence analyst forecasts range from herding of forecasts (Clement and Tse, 2005; Trueman, 1994); analysts' career concerns creating forecast bias (Bradley et al., 2022; Harford et al., 2019; Hong and Kubik, 2003); and exposure to terrorist shocks and natural disasters leading to greater pessimism in analyst earnings forecasts (Cuculiza et al., 2021; Kong et al., 2021).

ball results influence an analyst directly, and that these shocks could have economic significance. Secondly, what makes these shocks useful for our purposes is the precise and individual-specific nature of their sentimental impact. To illustrate with an example, consider an analyst who attended LSU, whose colleagues are Boston College graduates. In this case, the shock of LSU winning the championship is plausibly isolated to the focal analyst and doesn't affect their coworkers other than through spillovers. It is then possible to assess whether the shock of 'winning' influences the forecast behavior of the analyst's *peers*; i.e. the other analysts that work in their office.

To perform this analysis, we require data on analysts' college attendance. This information is not readily accessible in standard databases of analyst forecasts, such as IBES, prompting us to collect this information from a variety of sources. Using a procedure we describe in Section 2.2, we combine three different datasets containing analyst-level attributes (Bloomberg, CapitalIQ, and LinkedIn). Of the 57,749 individual analysts in IBES, we are able to link 7,481 analysts to schooling information. We are therefore working with a sample that is more than three times larger than that in Cohen et al. (2010). To the best of our knowledge, our data is the most comprehensive on analyst education currently assembled.

As more prominent individuals are naturally easier to identify in our education data, the matched sample consists of analysts covering more, larger (by asset size), and more profitable firms. These individuals furthermore release more forecasts and tend to be longer tenured than their peers. While this is not surprising, it does prompt us to restrict the control group to other matched analysts, alleviating concerns that our results are driven by ex-ante differences in the sample composition.

With this sample in place, we turn to estimating the treatment effects on individuals and the magnitude of spillovers on their colleagues. We find two significant results. Firstly, we document that the shock of a college winning the NCAA championship game does indeed influence the forecasts of an analyst who is a graduate. Our estimates suggest that 'winners' post forecasts roughly 0.12 standard deviations higher, significant at the 5% level.[5] Furthermore, we provide evidence

[5] In untabulated results, we fail to find evidence that these forecast responses result in lower accuracy of forecasts. Given that previous work has documented the endogenous relationship between forecasts and earnings, this result is perhaps unsurprising. We discuss this endogeneity, and find firm-level results consistent with them in Section 2.6.1 of our paper.

that this shock permeates to analysts who work at the same brokerage as these 'winners.' Our estimates suggest that, if 10% of an analyst's colleagues were 'winners,' a 'non-winner' analyst posts forecasts roughly 0.04 standard deviations higher, significant at the 5% level. Critically, our findings hold controlling for analyst-firm fixed effects. Including these fixed effects allows us to alleviate concerns surrounding selection effects in analyst-firm relationships, as the estimated impact of victory/loss is measured *relative to the analyst-firm history*, rather than across or between analysts and firms. We further control for a host of firm-level fundamentals following So (2013).

We assess the robustness of these results through several additional exercises. Firstly, adding a firm-analyst-month fixed effect ensures that our results are not driven by within-year variation of analyst forecasts. Given that the number of colleges that have won the NCAA championship game across our sample period (2000-2020) is relatively small (only 10 unique colleges in total), we furthermore perform a 'leave-one-out' exercise in which we systematically remove winning colleges to check whether any individual school is responsible for the effect. Reassuringly, we consistently find the same basic effect. Dropping small firms (below $200m in total assets), those in finance, insurance, and real estate sectors, as well as utilities, non-operating establishments, and industrial conglomerates doesn't materially change our results either.

We also assess whether the 'expectation' of victory/loss matters for the individual analyst result. It is natural to suppose that an unexpected loss, for instance, would be more upsetting than an expected one. We collect historical odds from betting markets for the NCAA finals in our sample, which are then interacted with the treatment variable in our difference-in-differences estimation. We find that whilst odds are unrelated to the impact for winners, their inclusion creates a statistically significant drop in forecasts for losers, with more unexpected losses inducing greater drops in subsequent forecasts that year. This result is therefore consistent with the idea that the 'expected' result matters for subsequent analyst sentiment.

Lastly, we conduct two placebo tests. In the first, we randomly assign a 'winner' indicator to analysts who we know were not in fact winners, followed by re-estimating the treatment effect. Repeating this exercise 2,000 times yields a distribution of the relevant coefficient that is centered around zero, with a mean and median of -.0005 and -0.0044 standard deviations respectively. 87.5% of the estimated coefficients are lower than in our initial analysis. In the second placebo test, we randomly assign a 'proportion' of winning analysts to a given brokerage in a given year that did not employ any 'winners,' followed by estimating our preferred regression specification. The proportions that we use for our this test are drawn from those observable in the data. The distribution of treatment effects after 2,000 samples is once again centered around zero. In total, 99.65% of the estimated placebo coefficients are lower than our original estimate.

Finally, to ensure that our brokerage spillover effect is driven by peer interaction and not proximity to a treated analyst's college town, we perform an exclusion exercise where we remove all US-based analysts operating outside of New York. New York has never had a college football team participating in the NCAA final, while accounting for 93.6% of all US-based analysts. Once again, we find a significant and positive spillover effect of winners to peers.

With these results in place, we proceed to investigate whether observable features of an analyst's peer group influence the nature of these spillovers. Given that organizational cultures can vary widely from company to company (Handy, 2007), it is probable that some brokerages are more likely to induce diffusion of the football sentiment shock to peer analysts than others.

To test this hypothesis, we estimate a brokerage-specific 'winner spillover' effect, which we will refer to as the 'Bro-ness.' The raw score we consequently refer to as 'Bro Score,' while the ranked measure we will call the 'Bro Rank.' We use the same specification as before, controlling for a firm-analyst fixed effect, a month-year fixed effect, and a host of controls.

Subsequently, we investigate whether these brokerage-specific effects correlate with observable features of the brokerages. Since men are nearly three times more likely to describe themselves as 'avid fans' of college football than women (37% of men vs. 14% of women),[6] we first assess whether brokerages that are more responsive to these kinds of sentiment shocks (i.e., have higher 'Bro Scores') differ in their share of female analysts. Consistent with our intuition, brokerages that respond more strongly to the sport sentiment shocks of peers have a significantly lower proportion

[6]This finding was reported in a recent (January 2023) poll of 2,201 American adults, published by Morning Consult, a global decision intelligence company.

of female analysts. Moving from the brokerage with the lowest response to the highest is associated with a fall in the total percentage of female analysts working at the brokerage of 6.24%. Given that the average female representation at a given brokerage across our sample is only 12.6%, this is a statistically and economically significant difference.

Another notable feature of sports culture is a stereotype of unpleasant and/or problematic male behavior. With this in mind, we collect two sets of measures of corporate 'misbehavior', with a view to assessing whether our 'Bro Score' correlates with these measures. The first is MSCI/KLD Scores. As prominent measures of firm-level ESG performance, these scores are typically used in the literature to assess the degree of corporate social responsibility in firms.[7] Our second measure is the corporate culture scores developed in Li et al. (2021). These scores, constructed using machine learning techniques applied to quarterly earnings call data, identify the strengths of firm culture along five dimensions: (1) Integrity, (2) Teamwork, (3) Innovation, (4) Respect, and (5) Quality. We find mixed evidence that 'Bro Scores' correlate with ESG scores, though the overall picture suggests a negative relationship. By contrast, using the corporate culture scores from Li et al. (2021), we find that 'Bro Scores' are associated with higher degrees of 'Teamwork' and 'Respect'.

As a final exercise, we assess the firm and stock market implications of our sentiment shocks. Since the treatment of analysts is plausibly exogenous to the business environment of the firms they cover, we construct an instrument based on the share of winners covering a given firm. This instrument is then used in a standard instrumental variable (IV) setting to establish a plausibly causal relationship between forecasts and firm/stock market level objects. We consider two outcome variables: (1) firm-level earnings-per-share (EPS), and (2) abnormal stock returns, computed using a Carhart (1997) four-factor model.

Consistent with Shore (2023), we document a near one-to-one response of firm-level earnings to a consensus forecast shock induced by an analyst 'winning.' Secondly, and again consistent with Shore (2023), there does not seem to be any evidence that the forecast shock moves the market: a one standard deviation increase in the consensus forecast, driven by the sentiment shock,

[7] See Gillan et al. (2021) and Hong and Shore (2022) for recent reviews of the Corporate Social Responsibility literature, where the use of MSCI/KLD scores is very widespread.

leads to moves in abnormal returns by between -0.15% to -0.83% depending on specification, not statistically different from zero. We can rule out that these results are due to insufficient power in the first stage (F-stat of 194.96 to 290.37). These findings suggest that the market is able to parse out arbitrary variation from forecasts. That said, these findings are to be interpreted with caution due to the small number of winners in our sample.

Our paper contributes to several strands of literature. Firstly, we add to the growing literature on the importance of social networks in economic decision-making (Bailey et al., 2018a,b, 2022). We also document the significance of analyst-specific sentiment shocks to their forecasting behavior (Cuculiza et al., 2021; Edmans et al., 2007; Kempf and Tsoutsoura, 2021; Kong et al., 2021). Finally, we contribute to the literature that identifies how firms actively responds to analyst forecasts by managing their earnings (Almeida et al., 2016; Bhojraj et al., 2009; Shore, 2023; Terry, 2015).

The remainder of this paper is organized as follows: Section 2.2 describes the data collection process and presents summary statistics of our final matched dataset. The empirical approach for identifying analyst and brokerage-level responses to the NCAA championship sentiment shock is outlined in Section 2.3. Section 2.4 presents our main results. Subsequently, Section 2.5 provides details on the analysis of the variation in brokerage-specific responses to sentiment shocks and the underlying characteristics of those brokerages. Firm and stock market responses to our sentiment forecast shocks are assessed in Section 2.6. Our results are discussed in Section 2.7. Ultimately, Section 2.8 concludes.

2.2 Data

The construction of our novel dataset is a major contribution of this paper. In this section, we outline the different data sources we draw from and how they are combined to arrive at the final sample. Earnings forecasts from IBES, which include the initial and (abbreviated) last name of an analyst as well as a brokerage identifier are the starting point for the sample construction. After unmasking the brokerage names we substantiate the dataset with novel data on education, employment, and gender from Bloomberg, CapitalIQ, and LinkedIn. This process leads us to identify the education of 7,481 analysts, a sample that is more than three times larger than that of Cohen et al. (2010). We finalize the dataset by adding firm-level controls from Compustat and CRSP.

2.2.1 Forecast Data

In line with the existing literature on analysts' forecasts, we obtain the earnings-per-share (EPS) estimates for the current and following fiscal year from IBES. We start from the universe of sell-side analysts in IBES which compromises a total of 67,300 analysts. Since our shocks are analyst-specific, we remove forecasts that are attributable to research teams and desks. We identify and then eliminate entries with either multiple names in the name field (8,192), or those that contain names of research desks (1,359).[8] This leaves us with 57,749 unique analysts for whom we have a last name and initial.[9]

A roadblock in identifying individual analysts and brokerages in the data is related to a change made to the dataset in October 2018 when the vendor began anonymizing analyst and brokerage identifiers. In contrast to earlier studies (Cohen et al., 2010), we therefore need to first unmask the brokerage identifiers, followed by linking analysts to our remaining data. We do so following the methodology proposed by Gibbons et al. (2020). For each entry in the recommendation table ('RECDET'), we obtain the name of the analyst ('ANALYST'), the individual's identifier ('AMASKCD'), and the identifier of the brokerage the analyst is working at ('ESTIMID'). Since the brokerage identifiers are derived from institutions' names, we manually link the 94 most prevalent ESTIMIDs (i.e. largest brokerages) to the brokerage names. Analysts working at these entities are responsible for 46% of all forecasts.

[8] A research desk usually carries the name of an industry/region or a (truncated) form of 'research department.'

[9] According to Fang and Hope (2020) a large number of observations in IBES are attributable to teams despite only the name of the lead analyst being listed in the database. To the extent that we don't already address this concern by excluding observations from analyst teams, this should only bias our results downwards.

2.2.2 Analyst-Level Attributes

Employment and educational background of analysts are obtained from three different sources: (1) Bloomberg, (2) CapitalIQ, and (3) LinkedIn. As the analyst names in IBES consist of the initial and last name, potential matches are identified based on an exact match of the (truncated) last name and an exact match of the initial. To zero in on valid matches, we require an overlap between the employers as derived from IBES and those reported in the outside dataset. Once we have obtained the (plausibly) complete name of an analyst, we use that information to specifically search for that individual across datasets.

It should be noted that neither of the three sources necessarily provides a complete record of an analyst's education.[10] In an attempt to alleviate these concerns, we construct plausibly complete educational histories for all of the matched analysts by drawing from data across all datasets. The names of all universities appearing in the datasets are disambiguated and assigned unique identifiers. We further obtain information on the degree and major earned from the institution and, where available, the years of attendance.

Of the 57,749 individual analysts in IBES we are able to link 13,307 (23.04%) to at least one of our sources for career and education information. Since not every entry contains information on educational attainment, we end up with schooling information for 7,481 analysts. We are ultimately working with a sample that is more than three times larger than that in Cohen et al. (2010).[11] To the best of our knowledge, our data is the most comprehensive on analyst education compiled to date.

Table 2.1 presents summary statistics of the education data. To give a sense of the analysts in our matched sample we construct a series of dummy variables that are commonly used in the corporate governance literature. The dummy variables in Panel A of Table 2.1 take on values of one if the analyst has obtained a Bachelors/Masters/Doctoral degree respectively and zero otherwise.

[10] Based on our investigations it is not uncommon for individuals to omit undergraduate education on LinkedIn, exclusively listing graduate degrees instead.

[11] Cohen et al. (2010) end up with a sample of 1,820 unique analysts for whom they gather educational data from zoominfo.com among others.

Table 2.1: Analyst Education

This table presents the summary statistics of the analyst dataset. Data from Bloomberg, CapitalIQ, and LinkedIn are disambiguated and cross-verified where an individual analyst is found in more than one source. We construct a series of variables to quantify the educational background of the individuals in our sample. Panel A presents the summary statistics of dummy variables that capture the attainment of a Bachelors, Masters, or Doctoral degree. The variables presented in Panel B are related to university attendance, capturing location (US and UK dummies) and Ivy league membership (IvyLeague). Lastly, Panel C presents the summary statistics for the analyst gender and year of birth (since we are working with a panel dataset, the year of birth is more meaningful than age).

Statistic	**N**	**Mean**	**SD**	**Min**	**Max**
Panel A: Degree					
Bachelors	6,351	0.774	0.419	0	1
Masters	6,351	0.599	0.490	0	1
PhD	6,351	0.054	0.225	0	1
Panel B: University					
US school	6,181	0.536	0.499	0	1
UK school	6,181	0.163	0.369	0	1
IvyLeague	6,114	0.136	0.342	0	1
Panel C: Gender & Age					
Female	6,336	0.137	0.344	0	1
Year of Birth	764	1968	9.521	1929	1994

As mentioned previously, while we observe educational information for all analysts in this sample, we don't necessarily have a full record of the degrees that were obtained. Consequently, only 77.4% of individuals have a Bachelors. While we would expect all analysts to have completed undergraduate education, there are two reasons why we don't observe the attainment of a B.A. or equivalent for about one quarter of them. Firstly, some analysts report the attendance of a university without stating the terminal degree that was obtained, which biases down this number. Secondly, there are analysts who omit their undergraduate education and exclusively list the completion of an advanced degree (most commonly an MBA).

While more than half (59.9%) of the analysts in our sample have completed a graduate degree, only about one in twenty hold a doctoral degree. With respect to the universities attended, Panel B reports summary statistics for a series of dummy variables that are designed to capture the location of their education. 53.6% of individuals obtained at least one of their degrees from a university in the US. Universities in the United Kingdom are the second most frequently attended (16.3%). Somewhat surprisingly, one in four of the analysts that obtained at least one of their degrees from a US institution did so at an Ivy League university. As a point of comparison, among US board members the share of Ivy League educated individuals is about 19%.

Lastly, Panel C of Table 2.1 presents information on the age and gender of the analysts. Consistent with previous studies, women are underrepresented in the finance industry, with only about one in six (13.7%) analysts being female. While the coverage of the gender variable is very good, we only observe the year of birth for a very small number of the matched analysts. We therefore only report this information for exhibition purposes.

While our matching approach does not favor specific analysts conditional on those analysts being in one of the three data sources, we do expect the selection of analysts into these datasets to be non-random. In Table 2.2 we therefore compare the matched to the unmatched analysts at the analyst-year. Unsurprisingly, there are stark differences between the two groups. As more prominent individuals have stronger incentives to have accounts with Bloomberg/CapitalIQ/LinkedIn, the matched sample consists of analysts covering a larger number of bigger and more profitable firms. These individuals are furthermore more productive and longer tenured than the individuals in the unmatched sample. While this is not surprising, it does prompt us to restrict the control group to other matched analysts to alleviate concerns that our results are driven by ex-ante differences in the sample composition.

As the treatment variable will ultimately depend on university attendance, Table 2.3 presents the ten most common schools in our sample, as well as those that won/lost an NCAA final since 2000 (the start of our sample). We rank schools by the total number of unique analyst-school observations. Unsurprisingly, a large number of analysts attended Columbia University and New York University, two of the largest feeder schools for jobs on Wall Street. Outside the US, the universities of Oxford and Cambridge are the most attended schools. Naturally, fewer individuals attended the large football schools at any point in their life.

Another strength of our data is that we are able to observe the office location of most of the

Table 2.2: Matched vs. Unmatched Analysts

This table presents a comparison between the matched and unmatched analysts. Our comparison is time-invariant, i.e., we collapse all time-series variables down to a single observation. As we rely on matching names and employment when identifying IBES analysts in our other sources (Bloomberg, CapitalIQ, LinkedIn), we are naturally biased towards those individuals that are more prominent, leading to the observable differences between the matched and unmatched parts of our sample.

	Matched			**Unmatched**			Δ
	N	Mean	S.D.	N	Mean	S.D.	
Individual	17,846	0.998	0.039	103,244	0.950	0.218	0.049***
Firms Covered (Total)	17,846	11.752	8.297	109,001	7.655	7.453	4.096***
Firms Covered (CRSP)	17,846	9.768	7.960	109,001	6.211	7.009	3.557***
Firms Covered (Comp)	17,846	9.757	7.956	109,001	6.205	7.003	3.552***
Forecasts (Total)	17,846	75.618	64.893	109,001	42.715	53.944	32.903***
Forecasts (Mean)	17,846	5.951	3.250	109,001	4.662	3.222	1.289***
Tenure	17,846	6.653	5.277	109,001	4.379	4.527	2.274***
Individual	17,846	0.998	0.039	103,244	0.950	0.218	0.049***
Size (Mean)	14,108	8.010	1.490	74,323	7.785	1.679	0.224***
EPS (Mean)	14,026	1.532	1.835	73,933	1.329	1.871	0.203***
NI (Mean)	14,111	657.625	973.640	74,338	556.761	965.117	100.863***
ROE (Mean)	14,055	0.006	0.084	73,611	0.003	0.090	0.003***

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

analysts in the matched sample throughout their careers. In Figure 2.1 we plot the geographical distribution of analysts, inversely weighted by the number of positions held by an analyst. Unsurprisingly, the large (western) financial centers dominate the sample.

Lastly, we compare the three groups, winners, losers, control, within the matched sample with respect to the industries they cover to ensure that such differences are not driving our results. We do so by computing the share of firms covered by each analyst in a given year, where industries are defined based on one-digit SIC codes. Table 2.4 presents the relative distribution in coverage for the observations in the three subsamples. Our findings are consistent with no meaningful difference in industry coverage.

Table 2.3: Most Attended Universities

This table presents the ten most attended schools in our sample of analyst education, as well as those that played in the NCAA Football Championship at any point in time between 2000 and 2022. For each university the number of different analysts that attended it at some point is computed. We disambiguate the raw university names and assign unique identifiers to each one. Observations therefore refer to a university (e.g. Columbia University) irrespective of the school attended (e.g. Columbia Business School).

Rank	**University/College**	**Analysts**
1	New York University	324
2	Columbia University in the City of New York	304
3	University of Pennsylvania	232
4	Harvard University	179
5	University of Oxford	177
6	University of Chicago	172
7	Cornell University	154
8	University of Cambridge	152
9	The London School of Economics and Political Science	140
10	University of Toronto	128
	⋮	
46	University of Texas	47
59	University of Southern California	41
67	University of Notre Dame	37
75	University of Florida	32
94	Virginia Tech	25
98	University of Miami	24
102	University of Georgia	24
119	The Ohio State University	20
180	University of Alabama	12
197	University of Oregon	11
230	Florida State University	9
251	Louisiana State University	8
278	Auburn University	7
283	Clemson University	7
286	University of Tennessee	6
322	University of Oklahoma	5
329	University of Nebraska	5

Figure 2.1: Distribution of Analysts

These figures plot the distribution of analysts by working place. Using detailed information on the office addresses of all individuals, we count the number of analysts in each location. We inversely weight by the number of different workplaces at the analyst level. Subsequently, we aggregate to the state level (for US locations, Panel b) and the country level (Panel a). Out of 9,512 analysts for whom we can observe locations, 5,183 are based in the US, of which 4,852 are working from an office in New York.

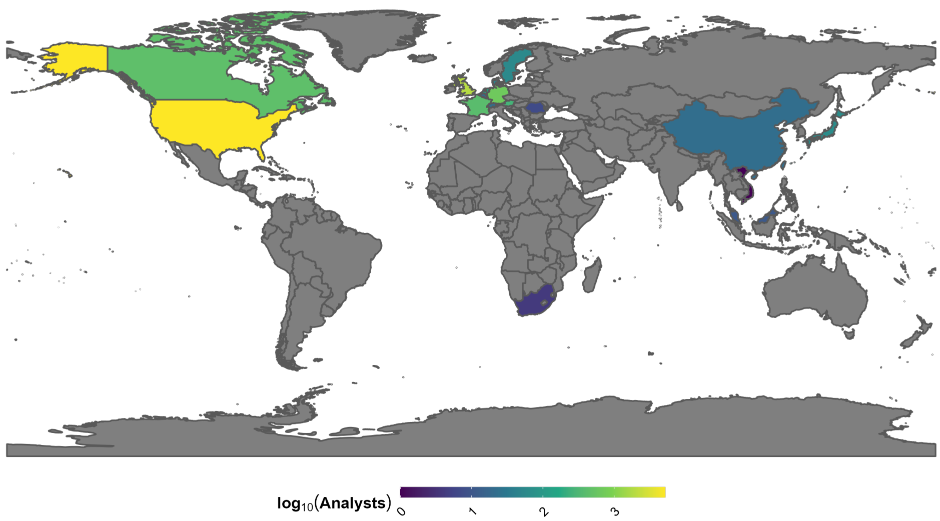

(a) World

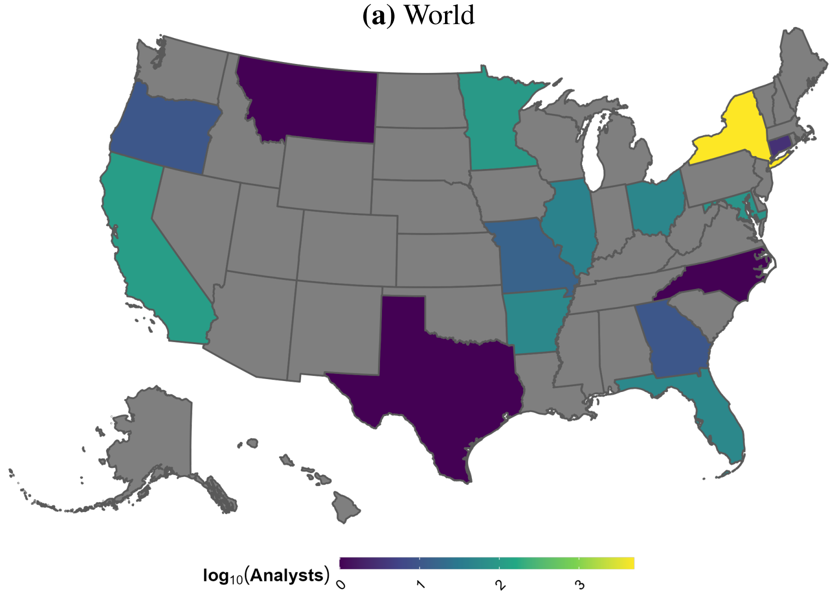

(b) United States

Table 2.4: Covariate Balance

This table presents a comparison in the industry coverage between the analysts in the treatment and control groups. For each analyst-year we compute industry coverage as the share of companies falling in each of the nine SIC industry ranges. The results when computing coverage based on total assets and market capitalization look similar.

	Winners (N = 255)			**Control (N = 50,500)**			**Losers (N = 240)**	
Variable	μ	σ	Δ_W	μ	σ	Δ_L	μ	σ
SIC 1	0.050	0.183	-0.042***	0.091	0.256	-0.006	0.085	0.253
SIC 2	0.269	0.407	0.105***	0.163	0.330	0.053**	0.217	0.376
SIC 3	0.292	0.395	0.068***	0.225	0.358	0.013	0.237	0.365
SIC 4	0.075	0.207	-0.031**	0.106	0.271	-0.008	0.099	0.249
SIC 5	0.085	0.213	-0.003	0.087	0.236	-0.017	0.070	0.206
SIC 6	0.051	0.186	-0.090***	0.141	0.325	-0.023	0.118	0.303
SIC 7	0.149	0.283	-0.002	0.151	0.297	-0.004	0.148	0.296
SIC 8	0.028	0.116	-0.005	0.034	0.133	-0.007	0.027	0.107
SIC 9	0.001	0.013	-0.000	0.002	0.022	-0.001***	0.000	0.001

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

2.2.3 Firm-Level Data

We draw company fundamentals from Compustat and stock prices from CRSP, linked to EPS forecasts from IBES. Firm variables that have been shown to influence analysts forecasts are used as controls. Guided by So (2013) we construct the following control variables from firm fundamentals: earnings-per-share when earnings are positive and zero otherwise, a binary variable indicating negative earnings, negative and positive accruals per share,[12] the percent change in total assets, a binary variable indicating zero dividends, dividends per share, the book-to-market ratio defined as book value scaled by market value of equity, and the end of fiscal year share price.

[12] Accruals equal the change in current assets (Compustat item ACT) plus the change in debt in current liabilities (Compustat item DCL) minus the change in cash and short-term investments (Compustat item CHE) and minus the change in current liabilities (Compustat item CLI).

Table 2.5: NCAA Football Finalists

This table presents the NCAA football finalists since 1999. For each year, the winning and losing team are presented. The finals take place in January in a 'neutral' location.

Year	Winner	Loser
2022	University of Georgia	University of Alabama
2021	University of Alabama	Ohio State University
2020	Louisiana State University	Clemson University
2019	Clemson University	University of Alabama
2018	University of Alabama	University of Georgia
2017	Clemson University	University of Alabama
2016	University of Alabama	Clemson University
2015	Ohio State University	University of Oregon
2014	Florida State University	Auburn University
2013	University of Alabama	Notre Dame University
2012	University of Alabama	Louisiana State University
2011	Auburn University	University of Oregon
2010	University of Alabama	University of Texas at Austin
2009	University of Florida	University of Oklahoma
2008	Louisiana State University	Ohio State University
2007	University of Florida	Ohio State University
2006	University of Texas at Austin	University of Southern California
2005	University of Southern California	University of Oklahoma
2004	Louisiana State University	University of Oklahoma
2003	Ohio State University	University of Miami
2002	University of Miami	University of Nebraska
2001	University of Oklahoma	Florida State University
2000	Florida State University	Virginia Tech
1999	University of Tennessee	Florida State University

2.2.4 NCAA Results

We obtain data on college sports from Nielsen's Gracenote and the NCAA.[13] For each year from 2000 to 2022 we collect the teams that competed in the NCAA football finals which generally take place in January of a given year. Table 2.5 contains a list of all finalists in our sample and whether they won or lost the game.

Analysts that went to the winning (losing) team's institution will be referred to as 'winners' ('losers') throughout the remainder of this paper. In each year we therefore separate our analyst sample into three groups: (i) 'winners,' (ii) 'losers,' and (iii) control. It's crucial to notice that the

[13]We access Gracenote through sports-reference.com, who supplement the raw data with additional information on teams. Data from the NCAA includes the location of all teams, their divisions, and a unique identifier that we use in matching datasets.

Table 2.6: Treatment/Control Groups

This table provides an overview of the number of analysts that are treated in a given year. In our empirical section we exploit the fact that an individual analyst covers a range of companies to achieve identification, alleviating concerns regarding the comparatively small number of treated individuals. Furthermore, since there are differences between matched and unmatched analysts, we restrict the control group to those analysts that were matched to their educational background.

Year	Winners	Losers	Control	Unmatched
2000	0	2	517	4,779
2002	5	0	624	4,980
2003	2	5	633	4,820
2005	10	0	686	4,129
2006	5	10	682	4,085
2007	0	3	716	4,124
2008	2	2	719	4,012
2010	0	4	761	3,935
2011	0	2	794	4,045
2012	0	1	799	4,065
2013	0	8	828	4,552
2014	3	0	895	5,432
2015	1	2	923	5,771
2018	0	2	856	5,064
2021	0	1	729	4,525

group membership changes year over year, i.e. an analyst that is classified as treated in t_0 can, and likely will, be in the control group in t_1. We present an overview of the relative sizes of the treatment and control groups in Table 2.6.

2.2.5 Final Sample and Summary Statistics

Starting from the forecasts that are made throughout the year, we generate a monthly time-series for each analyst-firm pair. Figure 2.2 presents an example of how our monthly time-series of forecasts is constructed, based on Google's parent company Alphabet Inc. (NASDAQ:GOOG). In transforming the data, we need to address two separate issues. The first, and more relevant one, is the timing of our exogenous event in relation to the fiscal year end of most US corporations. As the NCAA finals take place in January, the treatment coincides with the earnings season. When working with the forecasts for the earnings release closest in time, this leads to a natural truncation of the time-series.

To make this point clearer, Figure 2.2a depicts the dispersion of EPS forecasts made for Alphabet's 2020 fiscal year, which ended on December 31, 2020. As has been documented many times before, most forecast updates are clustered around the quarterly earnings releases. More importantly though, there are only very few forecasts made after the close of the fiscal year. This naturally renders the fiscal year 2020 ill-suited for testing the effect of a treatment in January 2021.

However, as most analysts release forecasts for several future periods, we instead construct our monthly time-series based on the forecasts for the 2021 fiscal year. As can be seen from Figure 2.2b, forecasts for FY2021 are available throughout the calendar years 2020 and 2021. This allows us to work with a significantly longer time-series. In cases where the fiscal year ends in a month other than December, we similarly work with the fiscal period giving us the most meaningful time-series of forecasts.

The second step in constructing the monthly time-series is a very intuitive one. As analysts update their forecasts at unpredictable frequencies, we construct our dataset from the most recent forecast by each analyst for a given company. Figure 2.2c depicts the time-series of EPS forecasts for Alphabet made by one of the analysts in our sample. The green triangles represent the time-series of forecasts pertaining to FY2020, underlining once more the necessity to work with the following fiscal period instead. The yellow diamonds are the time-series of the most recent EPS forecasts for FY2021 made by that analysts on the last day of a given month.

2.3 Empirical Strategy

In this section, we outline our empirical strategy for estimating analyst earnings forecast responses to college football sentiment shocks. Our approach involves a standard difference-in-differences framework, including analyst-firm fixed effects. Controlling for these fixed effects allows us to alleviate concerns surrounding selection effects in analyst-firm relationships. We begin by outlining how we estimate the analyst-level response to a direct shock of 'winning/losing.' We then discuss the construction of a brokerage specific treatment variable that measures the pro-

Figure 2.2: Forecast Revisions – Alphabet Inc. 2020/2021

Panel a depicts the forecasts of all sell-side analysts in IBES for Alphabet Inc. (NASDAQ:GOOG) made in 2020 and 2021 for the fiscal year ending December 31st, 2020. Panel b depicts the same information for the fiscal year ending December 31st, 2021. Panel c contains the interpolated forecast time-series for one analyst in the sample (`analys` 114392). Across panels the dashed line represents the 2020 fiscal year end.

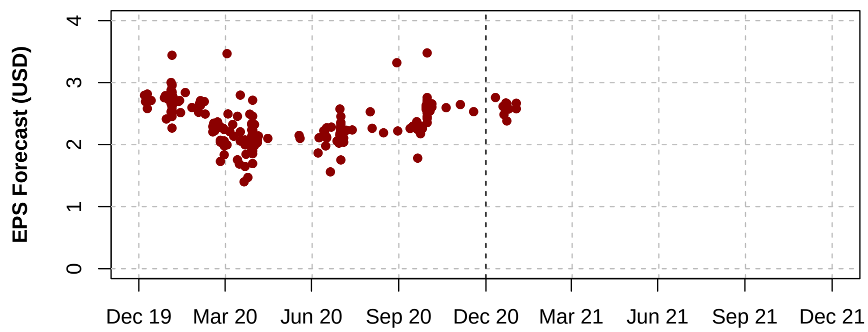

(a) FY2020 – All Analysts

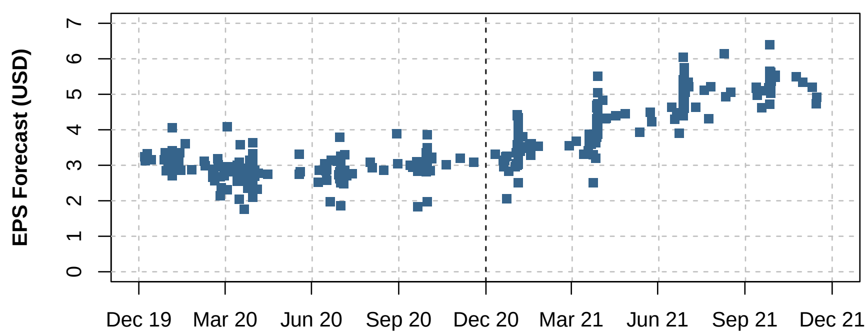

(b) FY2021 – All Analysts

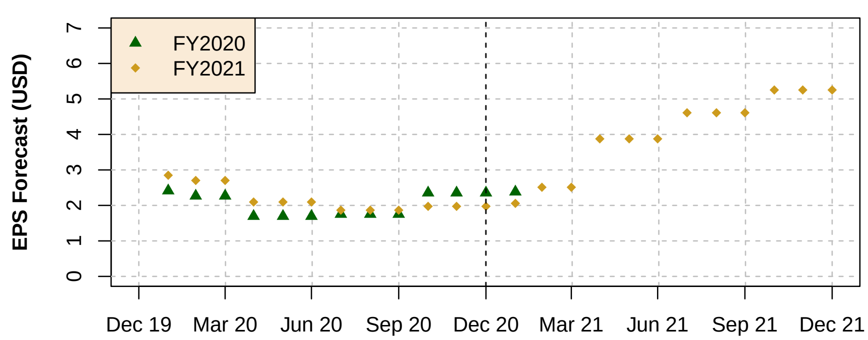

(c) Monthly Time-Series (Example)

portion of analysts affected by the college football shock. The construction of this variable allows us to measure within-brokerage spillovers, an advantage of our approach of identifying analyst-specific, rather than brokerage-wide shocks. Finally, we describe our approach for estimating firm and asset-pricing reactions to forecast shocks driven by the college football result.

2.3.1 College Sports

While sports events such as the NFL Super Bowl are known well beyond the borders of the US, the relevance of college sports is often not fully appreciated. Viewership of the 2023 NCAA football semi-finals and finals totaled more than 20mn per game. For comparison, the Super Bowl was viewed by 99mn individuals in 2022.

For our identification strategy to deliver a strong first stage it must be the case that having attended a college with a strong football culture leads the alumni to keep following the sport. Some of the schools in our sample are well known for investing heavily in their football programs, making it to the finals repeatedly over our sample period (see Table 2.5) and regularly having their graduates drafted into the NFL. Especially at these prominent schools (e.g., LSU), football is more than a pastime. Based on data compiled by Eren and Mocan (2018), the average attendance of Division I college football games was 45,000 in 2012 throughout the US. That said, it's not just the successful schools that draw fans from far and wide. In 1985 football dethroned baseball as the king of American sports. As most NFL franchises are located in large metro areas, there are vast parts of the country that are too far from the nearest venue to indulge in NFL games. As a natural consequence, college football had an equal rise in popularity.

College sports combine two core aspects of US culture, (1) alumni loyalty, and (2) enthusiasm for sports. It therefore doesn't come as a surprise that many graduates of schools with football programs develop a lifelong attachment to their *alma mater*'s team. Anecdotal evidence suggests that these ties remain strong even after having graduated and even moved to a different state.

2.3.2 Analyst Response

Our first goal is to show that winning/losing the NCAA Championship has a meaningful impact on analysts who attended the college that won/lost. To that end, we estimate the following regression specification:

$$\mathbb{E}_{i,j,t}[EPS_{j,y}] = \phi_{i,j} + \tau_t + \alpha_w w_{i,y} + \alpha_l l_{i,y} + \Gamma X_{j,y} + u_{i,j,t} \tag{2.1}$$

Where $\mathbb{E}_{i,j,t}[EPS_{j,y}]$ denotes the standardized earnings-per-share (EPS) forecast for firm j in year y, by analyst i at time (i.e. month-year) t.[14] $\phi_{i,j}$ is an analyst-firm fixed effect, τ_t is a month-year fixed effect, and $w_{i,y}/l_{i,y}$ is an indicator that takes a value of one if analyst i was a winner/loser in year y. X includes a host of covariates that have been shown to have predictive power for analysts' forecasts.[15] Our key parameters of interest are α_w and α_l. We interpret a positive/negative sign of these coefficients as evidence of a positive/negative shock to forecasts in the wake of victory/loss in the NCAA Championship College Football Final.

Identification is achieved through random variation in the identity of the school winning the NCAA final, even in the presence of endogenous matching between analysts and universities. Furthermore, by including analyst-firm fixed effects, we effectively compare an analyst's response relative to their own estimates in previous years. Nevertheless, the general concerns regarding DID estimated as outlined by Bertrand et al. (2004) remain to be addressed.

To assess whether the parallel trends assumption vital to DID estimators is violated, we estimate

[14] As earnings-per-share forecasts are often negative, we cannot take logs. To avoid any issues related to scaling, we standardize all forecasts by subtracting the firm-analyst level average forecast and dividing by the firm-analyst standard deviation. Thus our coefficients can be interpreted as standard deviations, rather than USD amounts.

[15] Following So (2013), who in turn bases their analysis on Fama and French (2006), we include the following lagged firm characteristics from year $y-1$: earnings-per-share when earnings are positive and zero otherwise, a binary variable indicating negative earnings, negative and positive accruals per share, where accruals equal the change in current assets (Compustat item ACT) plus the change in debt in current liabilities (Compustat item DCL) minus the change in cash and short-term investments (Compustat item CHE) and minus the change in current liabilities (Compustat item CLI), the percent change in total assets, a binary variable indicating zero dividends, dividends per share, book-to-market defined as book value scaled by market value of equity, and end of fiscal year share price.

a dynamic specification of the following form:

$$\mathbb{E}_{i,j,t}[EPS_{j,y}] = \phi_{i,j} + \tau_t + \beta_{-5} \sum_{s \le -5} D^s_{i,y} + \sum_{s \in [-4,-2]} \beta_s D^s_{i,y} + \sum_{s \in [0,10]} \beta_s D^s_{i,y} + \beta_{11} \sum_{s \ge 10} D^s_{i,y} + \epsilon_{i,j,t} \quad (2.2)$$

Going from (2.1) to (2.2), we replace the static treatment variables ($w_{i,y}$ and $l_{i,y}$) with a series of indicator variables $D^s_{i,y}$ for all months s. These indicators take values of one if the analyst was a winner/loser in year y and if the forecast is s months from the treatment event (i.e. the championship game). To allow for asymmetry in the contribution of the control variables we estimate separate specifications looking only at winners/losers, as well as including both treatment dummies simultaneously.

2.3.3 Network Response

The central goal of our paper is to identify how shocks to analysts can spread to their social network. To this end, we consider one salient network to the analyst: their workplace. Specifically, we investigate whether working as an analyst at the same brokerage as a 'winner'/'loser' has any meaningful effect on the forecasts of a 'non-winner'/'non-loser.'

To assess spillovers between analysts within the same brokerage, we re-estimate the specification in (2.1) and (2.2), replacing the treatment variable with the within-brokerage proportion of 'winners'/'losers.' This altered specification reflects the notion that we expect the strength of peer effects to be increasing in the number of coworkers affected by the event. Let $\widehat{W}_{i,y}$ denote the proportion of 'winners' working at brokerage i in year y, and let $\widehat{L}_{i,y}$ denote the proportion of 'losers' working at brokerage i in year y. Then, the specification we estimate takes the following form:

$$\mathbb{E}_{i,j,t}[EPS_{j,y}] = \phi_{i,j} + \tau_t + \alpha_w w_{i,y} + \alpha_l l_{i,y} + \beta_w \widehat{W}_{i,y} + \beta_l \widehat{L}_{i,y} + \Gamma X_{j,y} + u_{i,j,t} \quad (2.3)$$

Where $\mathbb{E}_{i,j,t}[EPS_{j,y}]$ denotes the forecast of analyst i for firm j at time period (month-year) t, $w_{i,y}/l_{i,y}$ are indicator variables taking a value of one if analyst i is a winner/loser, and $X_{j,y}$ is a vector of firm-year level controls.

We interpret the coefficient of β_w and β_l in (2.3) as capturing a 'spillover' effect of winners/losers in the analyst workplace. If $\beta_w > 0$ for instance, we interpret this as evidence that the positive shock experienced by a 'winning' colleague spreads to their colleagues forecasting behavior.

2.4 Results

In this section we detail our main results. We begin by showing that an analyst's forecasts do indeed respond to the personal shock of 'winning' the NCAA National Championship game. We then show that this shock spreads to analysts that work alongside 'winners' at the same brokerage. Finally, we discuss our robustness checks.

2.4.1 Analyst Specific

We start by estimating (2.1) using the monthly time-series data we constructed as outlined in Section 2.2. The coefficient estimates are reported in Table 2.7. To assess the dynamics of this effect, we estimate the dynamic specification in (2.2) and plot the coefficients in Figure 2.3.

We find that 'winners' post forecasts that are roughly 0.12 standard deviations higher in the wake of the shock of winning. Note that this coefficient is estimated in the presence of a firm-analyst fixed effect, and hence reflects a change in the analyst's forecasts relative to their forecasts of the same firm in years when they are not winners. We find a negative point estimate for losers, though without statistical significance.

In untabulated results, we estimate the effect of winning/losing on an analyst's forecast accuracy. We fail to find evidence that the forecasts of treated individuals are systematically worse than those of their peers. While this may seem curious at first, this finding is consistent with Shore (2023) as well as our analogous findings reported in Section 2.6.1. Since firms respond to changes in the EPS target, there exists an endogeneity between forecasts and subsequent earnings. We discuss this endogeneity in more detail in Section 2.6.1.

As is visible in Figure 2.3 above, there is a clear parallel trend between 'winners' and 'losers,'

Table 2.7: Static Estimation

This table presents the results of estimating equation (2.1). Here $w_{i,y}/l_{i,y}$ are indicators taking a value of one if analyst i was a 'winner'/'loser' in the NCAA Championship Final in year y. The controls are described in Section 2.3 and follow the specification in So (2013) for constructing characteristic forecasts. In all cases we control for a firm-analyst fixed effect and a 'month-year' fixed effect.

Dependent Variable:		$\mathbb{E}_{i,j,t}[EPS_{j,y}]$	
Model:	(1)	(2)	(3)
Variables			
$w_{i,y}$	0.1229***		0.1212***
	(0.0442)		(0.0465)
$l_{i,y}$		-0.0296	-0.0100
		(0.0526)	(0.0548)
Controls	Yes	Yes	Yes
Fixed-effects			
Firm-Analyst	Yes	Yes	Yes
Month-Year	Yes	Yes	Yes
Fit statistics			
Observations	3,072,433	3,072,433	3,072,433
R^2	0.94241	0.94241	0.94241

Clustered (Firm-Analyst) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

which adds credibility to our identification strategy. Two additional features of Figure 2.3 are worth commenting on: the gradual widening of the forecasts, rather than immediate jumps, and the persistence of the shock over time.

With respect to the former, it is worth noting, as discussed in Section 2.2, that forecasts occur continuously throughout the year, though infrequently. Across the whole sample, the median number of forecasts per analyst-firm-year is 4, with a standard deviation of 2.8. Across 'winners,' the median number of forecasts per analyst per firm-year similarly is 4, with a standard deviation of 2.3. As such, it seems reasonable that some time may pass before 'winning' analysts actually post forecasts.

To test this possibility, we identify the proportion of winning analysts whose first forecast comes after the month of the football game (January, or $t = 0$). We find that 64.3% of winning

Figure 2.3: Forecast – Dynamic

This figure plots the coefficients and 95% confidence intervals from our preferred difference-in-differences specification. We estimate the following regression $\mathbb{E}_{i,j,t}[EPS_{j,y}] = \phi_{i,j} + \tau_t + \beta_{-5}\sum_{s\leq -5} D^s_{i,y} + \sum_{s\in[-4,-2]}\beta_s D^s_{i,y} + \sum_{s\in[0,10]}\beta_s D^s_{i,y} + \beta_{11}\sum_{s\geq 10} D^s_{i,y} + \epsilon_{i,j,t}$ and extract the estimates $\widehat{\beta}_1$, which we then plot for two different definitions of treatment (a) 'winners' and (b) 'losers.' Standard errors are clustered at the analyst-level.

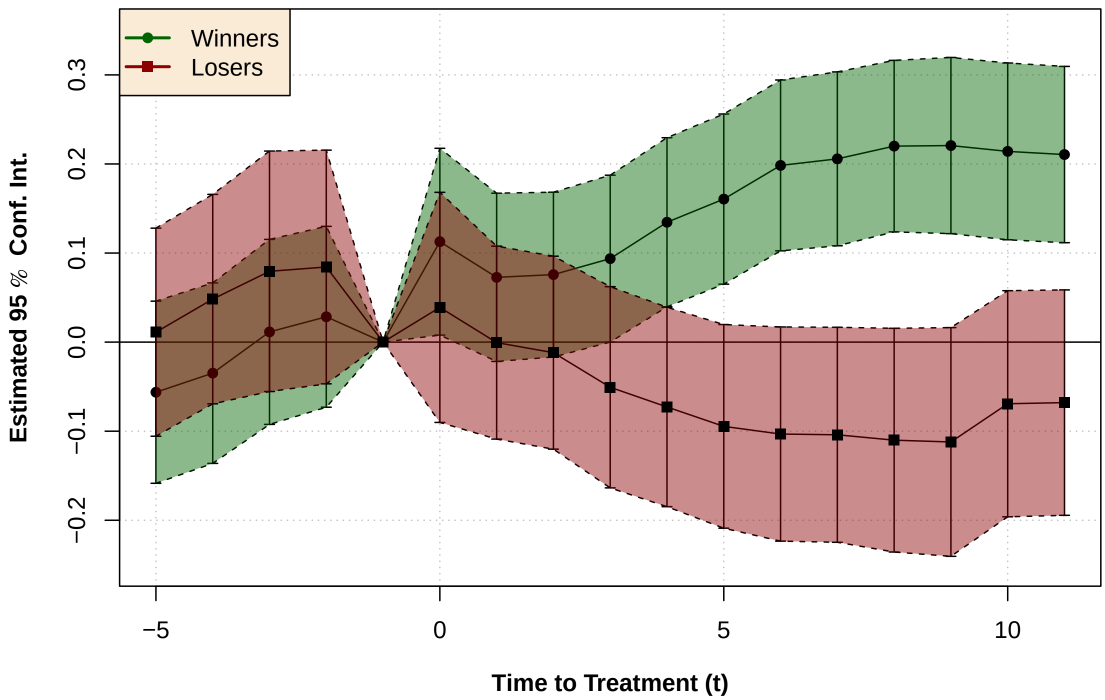

analysts post their first forecast at least one month after the game ($t = 1$), 32.4% at least two months after, 26.5% at least three months after, and 15.8% at least four months after. A sizeable 13.2% of winning analysts do not post their first novel forecast for six months after the championship game.

Note that this feature of the forecast data also goes some way towards explaining the persistence of the effect, given that the impact of the college football shock may not appear until late in the year. That said, a natural and valid concern is to question the plausibility of a college football game influencing mood for such a lengthy period that it would affect forecasts some six months after the fact. Here we defer to previous work that illustrates just how seriously US adults treat college football (Eren and Mocan, 2018).

We also check whether winners are more likely to revise negatively than non-winners, and if they do revise negatively, then whether winners do so more aggressively than the control group. Given that we observe no meaningful revision in the wake of the positive forecast shock, we should find that winners are no more likely to revise negatively, and that their negative revisions are no larger than for non-winners. We find that whilst winners post negative revisions roughly 1.03 times more often than positive revisions, this is slightly less than for non-winners, who post negative revisions 1.08 times more often than positive revisions. Similarly, when a revision is downward, the average for winners is -0.37 standard deviations, whereas for non-winners the average revision is -0.44 standard deviations. Thus these mean comparisons are consistent with our finding of a persistent effect on forecasts stemming from the football shock.

2.4.2 Brokerage Spillovers

Having established that treated analysts respond in the expected direction, we move on to answering the question of whether there are spillovers between analysts in the same brokerage. We do so by adding the variables $\widehat{W}_{i,y}$ and $\widehat{L}_{i,y}$, that measure the proportion of analysts working in the same brokerage as analyst i who were 'winners' and 'losers' respectively.

Out of a total sample size of 469 brokerages, we find 34 (33) brokerages that have non-zero values for $\widehat{W}_{i,y}$ ($\widehat{L}_{i,y}$) in at least one year. In total, around 22.3% of our brokerage-year sample contains a non-zero value for $\widehat{W}_{i,y}$. Around 8.1% of our forecast observations contain non-zero values for $\widehat{W}_{i,y}$, and roughly 7.7% for $\widehat{L}_{i,y}$. Within the forecast observations with non-zero values, the average value of $\widehat{W}_{i,y}$ is 3.8%, with a standard deviation of 4.6%, and the average value of $\widehat{L}_{i,y}$ is 4.0% with a standard deviation of 5.6%.

The coefficient estimates from running regression (2.3) are reported in Table 2.8. We find that the presence of winners in the workplace has a statistically significant positive effect on the forecasts of their peers. While optimism is contagious, with coworkers raising their forecasts by 0.04 standard deviations for every 10% of their colleagues who are winners, the losers' pessimism doesn't spread within the brokerage.

As before, we separately assess the dynamics of this effect by estimating (2.2) with the treat-

Table 2.8: Contagion

This table reports coefficient estimates from the following regression $\mathbb{E}_{i,j,t}[EPS_{j,y}] = \phi_{i,j} + \tau_t + \alpha_w w_{i,y} + \alpha_l l_{i,y} + \beta_w \widehat{W}_{i,y} + \beta_l \widehat{L}_{i,y} + \Gamma X_{j,y} + u_{i,j,t}$. $w_{i,y}$ ($l_{i,y}$) is a dummy that takes on a value of one for an analyst that is a winner (loser) of the college championship game. $\widehat{W}_{i,y}$ ($\widehat{L}_{i,y}$) represents the proportion of winners (losers) that analyst i works alongside in year y. Standard errors are clustered at the firm-analyst-level.

Dependent Variable:		$\mathbb{E}_{i,j,t}[EPS_{j,y}]$	
Model:	(1)	(2)	(3)
Variables			
$w_{i,y}$	0.0385		0.0368
	(0.0399)		(0.0422)
$\hat{W}_{i,y}$	0.4037**		0.4115**
	(0.1624)		(0.1636)
$l_{i,y}$		-0.0223	-0.0133
		(0.0564)	(0.0585)
$\hat{L}_{i,y}$		0.0804	0.1004
		(0.2023)	(0.2040)
Controls	Yes	Yes	Yes
Fixed-effects			
Firm-Analyst	Yes	Yes	Yes
Month-Year	Yes	Yes	Yes
Fit statistics			
Observations	2,931,311	2,931,311	2,931,311
R^2	0.93625	0.93625	0.93625

Clustered (Firm-Analyst) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

ment variable defined as above. The point estimates and confidence intervals are depicted in Figure 2.4. We observe a similar pattern as outlined in Section 2.4.1.

2.4.3 Robustness

To check for robustness, we perform five separate exercises. In the first, we control for a firm-analyst-*month* fixed effect in our first specification, seeking to alleviate concerns that our estimates are picking up some time variation in the strength of analyst forecasts. Secondly, we test whether the treatment effect varies with the expected outcome of the game by drawing on betting odds prior to the start of the game. Subsequently, we perform a 'leave-one-out' procedure wherein

Figure 2.4: Brokerage Effects

This figure presents coefficient estimates of the event study regression outlined in equation (2.2), with the treatment variables defined as $\widehat{W}_{i,y}$ for 'winners' of the NCAA Championship game, and $\widehat{L}_{i,y}$ for 'losers.' These treatment variables measure the proportion of analysts working in the same brokerage as analyst i who were winners and losers respectively. The coefficient should be interpreted as the impact on an analyst's forecast if 100% of their colleagues were 'winners.' We cluster standard errors at the firm-analyst level.

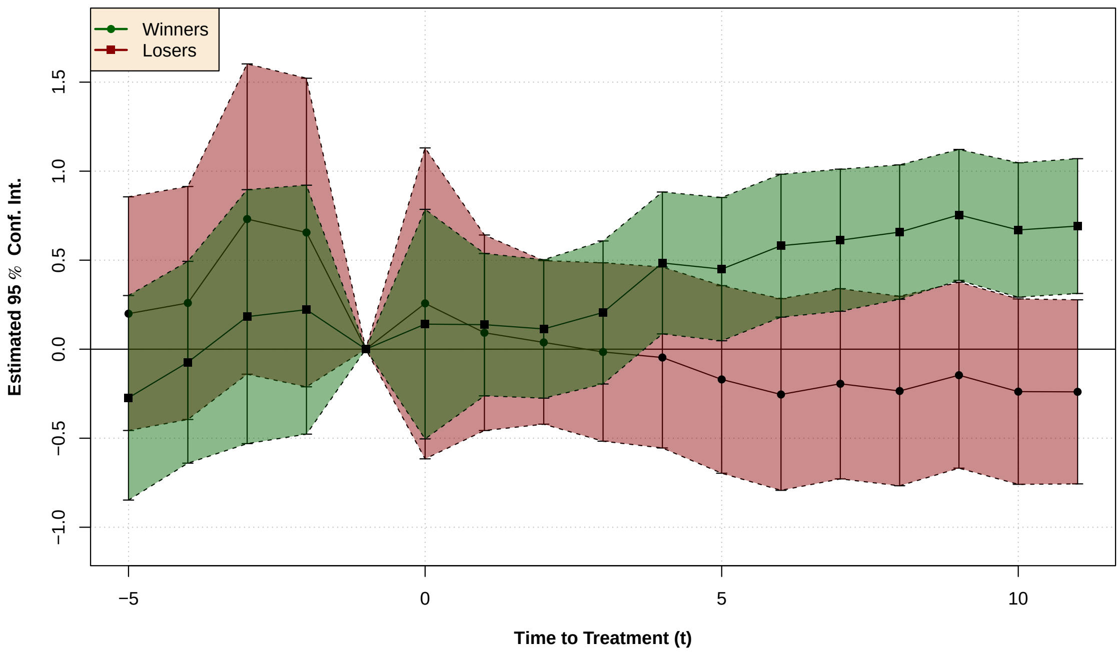

we systematically remove each winning school from our regression. It is our aim to show that no individual school is driving the headline result. In the fourth, we remove US-based analysts working outside of New York. Doing so allows us to eliminate the possibility that analysts who work in close proximity to colleges that win the NCAA championship are influenced by their local surroundings rather than their peers.[16] Finally, we conduct a standard placebo test.

Firm-Analyst-Month Fixed Effect

It is possible that analysts post more or less optimistic forecasts in certain months of the year. To ensure that such time variation is not driving our results, we control for a 'firm-analyst-month'

[16]Whilst 93.6% of analysts work in New York, no New York based college football team has ever appeared in the NCAA Championship Final game. See Figure 2.1 and Table 2.5 for more details.

Figure 2.5: Robustness – Firm-Analyst-Month Fixed Effect

This figure presents coefficient estimates of the event study regression outlined in equation (2.2), controlling for a firm-analyst-month fixed effect.

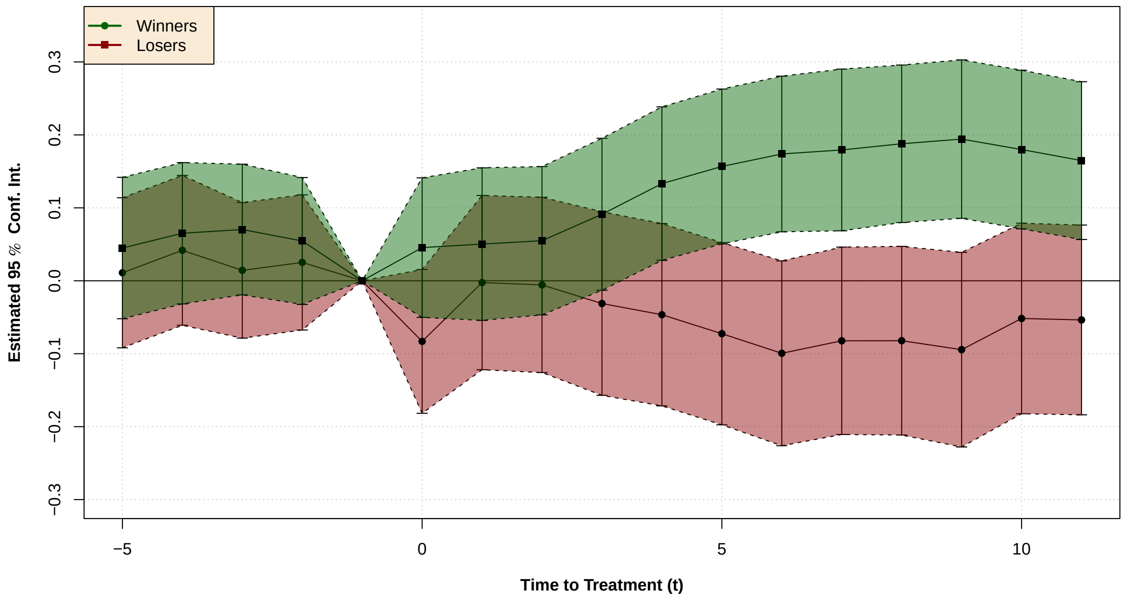

fixed effect. We are thus comparing an analyst's forecasts to the forecasts they posted for the same firm in the same calendar month (e.g. September), rather than across the entire year. Our specification remains otherwise unchanged. The results can be found in Figure 2.5; controlling for this time variation does not significantly affect our point estimates, nor the statistical significance of our results.

Unexpected Game Outcomes

One would expect that the size of the treatment effect varies depending on whether the game outcome was 'expected' or not. To establish a notion of 'close' games, we draw on odds from betting markets prior to the NCAA Championship game.[17] Odds are generally reported relative to a fair \$100 bet. If betting on Alabama winning over LSU currently has odds $-\$X$, a bet of $\$100+X$ will result in a profit of \$100 in case of payout. Negative odds can therefore be interpreted as the corresponding team being expected to win. In our sample, the odds for the winning team range

[17] We obtain the odds on the day of the game from *The Lines*.

Table 2.9: Robustness – Unexpected Wins/Losses

This table reports coefficient estimates from the following regression, $\mathbb{E}_{i,j,t}[EPS_{j,y}] = \phi_{i,j} + \tau_t + \alpha_w w_{i,y} + \alpha_w^o w_{i,y} \times odds_c^w + \alpha_l l_{i,y} + \alpha_l^o l_{i,y} \times odds_c^l + \beta_w \widehat{W}_{i,y} + \beta_l \widehat{L}_{i,y} + \Gamma X_{j,y} + u_{i,j,t}$. $w_{i,y}$ ($l_{i,y}$) is a dummy that takes on a value of one for an analyst that is a winner (loser) of the college championship game. $odds_c^x$ is the odds of college c attended by analyst i either winning ($x = w$) or losing ($x = l$) the NCAA final. Standard errors are clustered at the firm-analyst-month-level.

Dependent Variable:		$\mathbb{E}_{i,j,t}[EPS_{j,y}]$	
Model:	(1)	(2)	(3)
Variables			
Win	0.1154***		0.1035***
	(0.0124)		(0.0132)
Win × Winner Odds	-0.0018		0.0002
	(0.0020)		(0.0020)
Lose		-0.0645***	-0.0447***
		(0.0150)	(0.0158)
Lose × Loser Odds		0.0157***	0.0138***
		(0.0023)	(0.0024)
Controls	Yes	Yes	Yes
Fixed-effects			
Firm-Analyst-Month	Yes	Yes	Yes
Month-Year	Yes	Yes	Yes
Controls			
	Yes	Yes	Yes
Fit statistics			
Observations	3,072,433	3,072,433	3,072,433
R^2	0.98542	0.98542	0.98542

Clustered (Firm-Analyst-Month) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

from -10 (Alabama in 2013) to 11 (Ohio State in 2003), with a mean of -0.17. Naturally, the odds for the losing team are derived by multiplying those of the winner by negative one. Our findings are documented in Table 2.9.

Our results suggests that while winners are seemingly unaffected by the odds of their victory, losers react more strongly if their team was expected to win (i.e. the odds for a bet on their team were negative). Whilst in our static estimation, detailed in Table 2.7, we fail to find a statistically significant impact for losers, we now find a highly statistically significant and negative effect on subsequent forecasts, with more unlikely losses (i.e. negative loser odds) deepening the negative effect, and vice versa.

Leave-One-Out Estimation

One concern with our identification strategy is the relatively limited number of different schools that made it to the NCAA finals in our sample period (see Table 2.5). While the concern that individual schools are driving the results does not invalidate our approach as such, we strive to demonstrate that the treatment effect is not due to a single school. To alleviate such concerns, we return to our dynamic regressions specification in (2.2), dropping individual universities from the sample. Figure 2.6 presents the dynamic difference-in-differences plots for the ten samples when removing one of the finalists from Table 2.5.

As can be seen from Figure 2.6, the dynamic effect does not change substantially, in size and significance, when removing any of the focal schools. We therefore conclude that the results aren't driven by any single college.

Excluding US-based Analysts Outside of New York

To alleviate concerns that our results are driven by brokerages located in close proximity to colleges that win/lose the NCAA final, we exclude all US-based analysts outside of New York. In this respect, we limit the possibility that the local atmosphere is driving the forecast results.

Table 2.10 presents our findings from a simple difference-in-differences estimation where we once again find a positive and statistically significant brokerage effect. We also report the event study plot in Figure 2.7, where the same basic pattern emerges.

Placebo Test

We run two placebo tests designed to shed light on the plausibility of our two results, i.e. that 'winners' react to the NCAA Championship game, and that this effect spreads to their colleagues. In the first, we randomly assign the indicator $w_{i,y}$ to 'non-winners,' and set $w_{i,y}$ equal to zero for known 'winners.' We maintain the same proportion of 'winners' to 'non-winners' in our placebo

Figure 2.6: Robustness – Leave-One-Out

We assess the robustness of our results by re-estimating the dynamic effects when leaving out all analysts associated with one of the schools that won/lost the NCAA final. The panels in this figure correspond to the analogues of Figure 2.3 but with all observations related to the focal college removed from the sample.

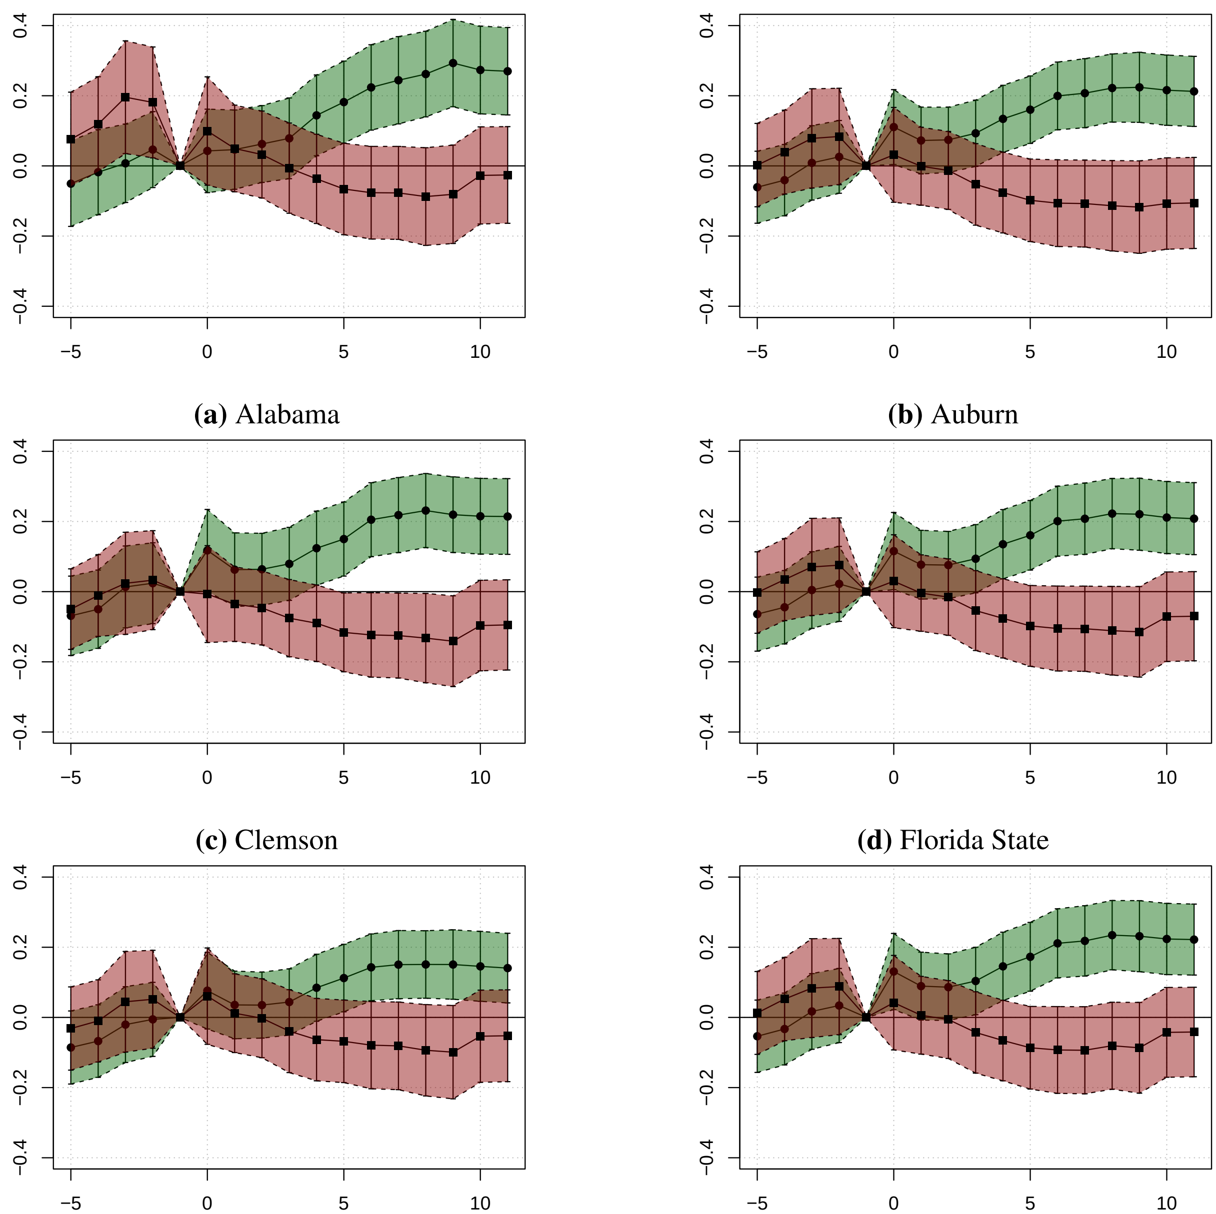

(a) Alabama

(b) Auburn

(c) Clemson

(d) Florida State

(e) LSU

(f) Miami

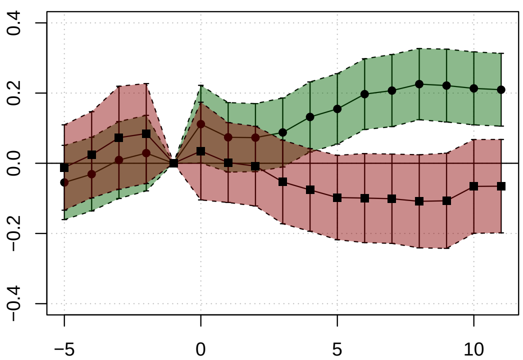

(g) Ohio State

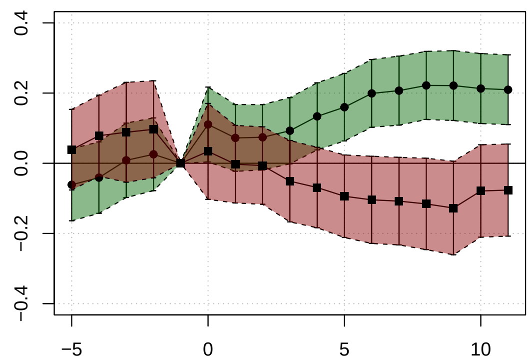

(h) Oklahoma

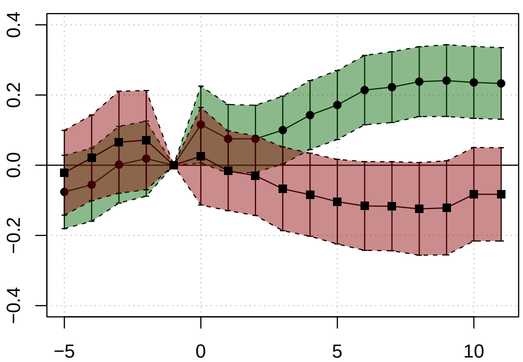

(i) USC

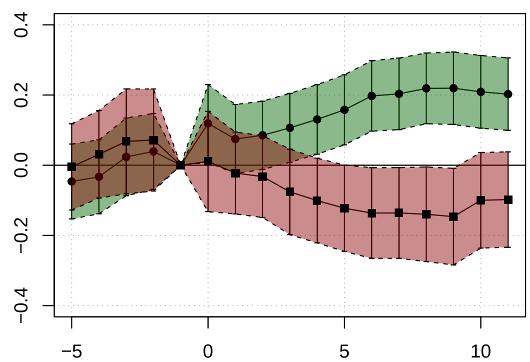

(j) Texas

Table 2.10: Robustness – NY and International Analysts

This table reports coefficient estimates from the following regression, $\mathbb{E}_{i,j,t}[EPS_{j,y}] = \phi_{i,j} + \tau_t + \alpha_w w_{i,y} + \alpha_l l_{i,y} + \beta_w \widehat{W}_{i,y} + \beta_l \widehat{L}_{i,y} + \Gamma X_{j,y} + u_{i,j,t}$, excluding from the sample all US-based analysts operating outside of New York. $w_{i,y}$ ($l_{i,y}$) is a dummy that takes on a value of one for an analyst that is a winner (loser) of the college championship game. $\widehat{W}_{i,y}$ ($\widehat{L}_{i,y}$) represents the proportion of winners (losers) that analyst i works alongside in year y. Standard errors are clustered at the firm-analyst-level.

Dependent Variable:		$\mathbb{E}_{i,j,y}[EPS_{j,y}]$	
Model:	(1)	(2)	(3)
Variables			
$w_{i,y}$	0.0504***		0.0486***
	(0.0122)		(0.0128)
$\hat{W}_{i,y}$	0.3587***		0.3640***
	(0.0525)		(0.0529)
$l_{i,y}$		-0.0245	-0.0132
		(0.0169)	(0.0174)
$\hat{L}_{i,y}$		0.0514	0.0690
		(0.0643)	(0.0648)
Controls	Yes	Yes	Yes
Fixed-effects			
Ticker-Analyst-Month	Yes	Yes	Yes
Month-Year	Yes	Yes	Yes
Fit statistics			
Observations	2,509,621	2,509,621	2,509,621
R^2	0.98811	0.98811	0.98811

Clustered (Ticker-Analyst-Month) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

test. In the second, we randomly assign a proportion of 'winners' in a given brokerage-year, $\widehat{W}_{i,y}$, to brokerages that we know had no 'winners' in that year, and set $\widehat{W}_{i,y}$ equal to zero for brokerages with known non-zero proportions. When randomly assigning proportions of 'winners,' we resample $\widehat{W}_{i,y}$ from the empirical distribution. In both cases, we perform these tests 2,000 times and report the distribution of coefficients.

Figure 2.8 presents the results of our first exercise. The distribution of estimates is centered around zero, with 87.5% of the estimated coefficients lying to the left of our main analysis estimate. It is worth noting that, across our entire sample, we have very few 'winners' relative to the size of the control group (see Table 2.6). This may explain why the distribution of placebo estimates is

Figure 2.7: Placebo Test – NY and International Analysts

This figure presents coefficient estimates from the following regression $\mathbb{E}_{i,j,t}[EPS_{j,y}] = \phi_{i,j} + \tau_t + \alpha_w w_{i,y} + \alpha_l l_{i,y} + \beta_w \widehat{W}_{i,y} + \beta_l \widehat{L}_{i,y} + \Gamma X_{j,y} + u_{i,j,t}$, excluding from the sample all US-based analysts operating outside of New York. $w_{i,y}$ ($l_{i,y}$) is a dummy that takes on a value of one for an analyst that is a winner (loser) of the college championship game. $\widehat{W}_{i,y}$ ($\widehat{L}_{i,y}$) represents the proportion of winners (losers) that analyst i works alongside in year y. Standard errors are clustered at the firm-analyst-level.

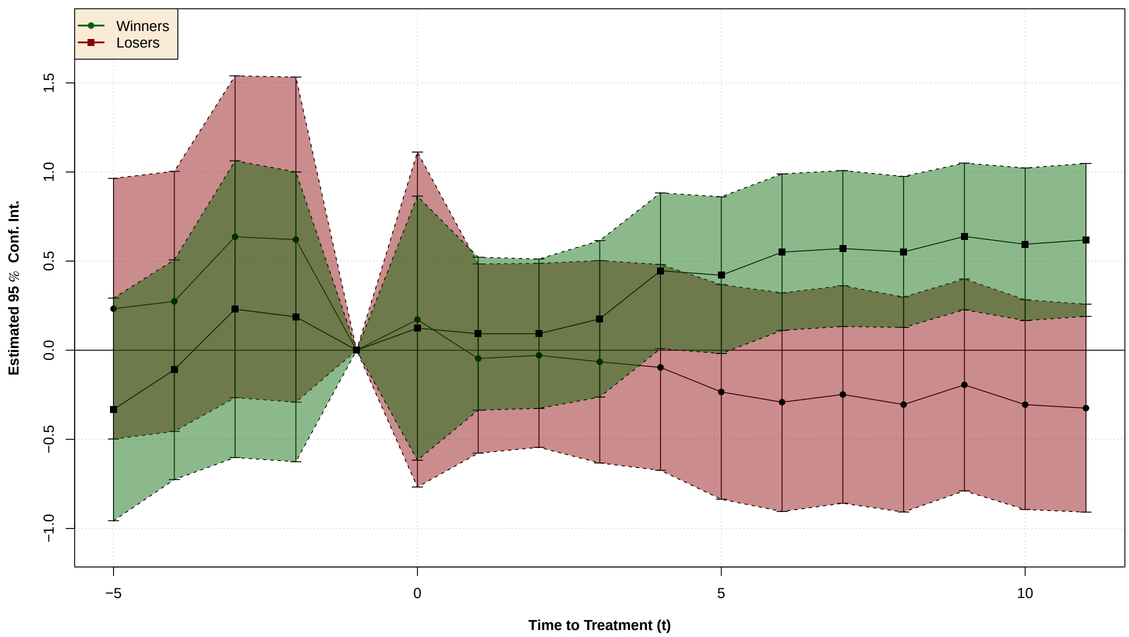

fairly broad, and why there are not an insignificant number of estimates that exceed our results in the non-placebo case.

Figure 2.9 shows the results of our second exercise. The distribution of estimates is centered around zero, with 99.65% of the estimated coefficients lying to the left of our main analysis estimate. In comparison to our first case, we have many more brokerage years affected by the variable $\widehat{W}_{i,y}$, compared to $w_{i,y}$. This feature of our second exercise offers one explanation for why the distribution over placebo coefficients is significantly tighter around zero, as compared to the first.

2.5 Within-Brokerage Spillovers and 'Bro-ness'

Having established that the college football shock influences winning/losing analysts, and that the shock of winning appears to spread to peers who work alongside winners, we now turn to

Figure 2.8: Placebo Test – Effect on 'Winners'

This figure presents coefficient estimates of our first placebo test. In that exercise, we randomly assign a winner indicator, $w_{i,y}$, to known non-winners, whilst setting the same indicator to zero for all known winners. The specification is then the same as in our main analysis: $\mathbb{E}_{i,j,t}[EPS_{j,y}] = \phi_{i,j} + \tau_t + \alpha_w w_{i,y} + \alpha_l l_{i,y} + \Gamma X_{j,y} + u_{i,j,t}$. We repeat this 2,000 times and report the distribution of our coefficient of interest, α_w, in the figure. We find a distribution of coefficients centered around zero, with a mean of -0.0005, and a median value of -0.0044. We include our coefficient estimate from our main analysis (0.1212) as the dotted line to the right of 0. Roughly 87.5% of our placebo test coefficient estimates are lower than our main estimate.

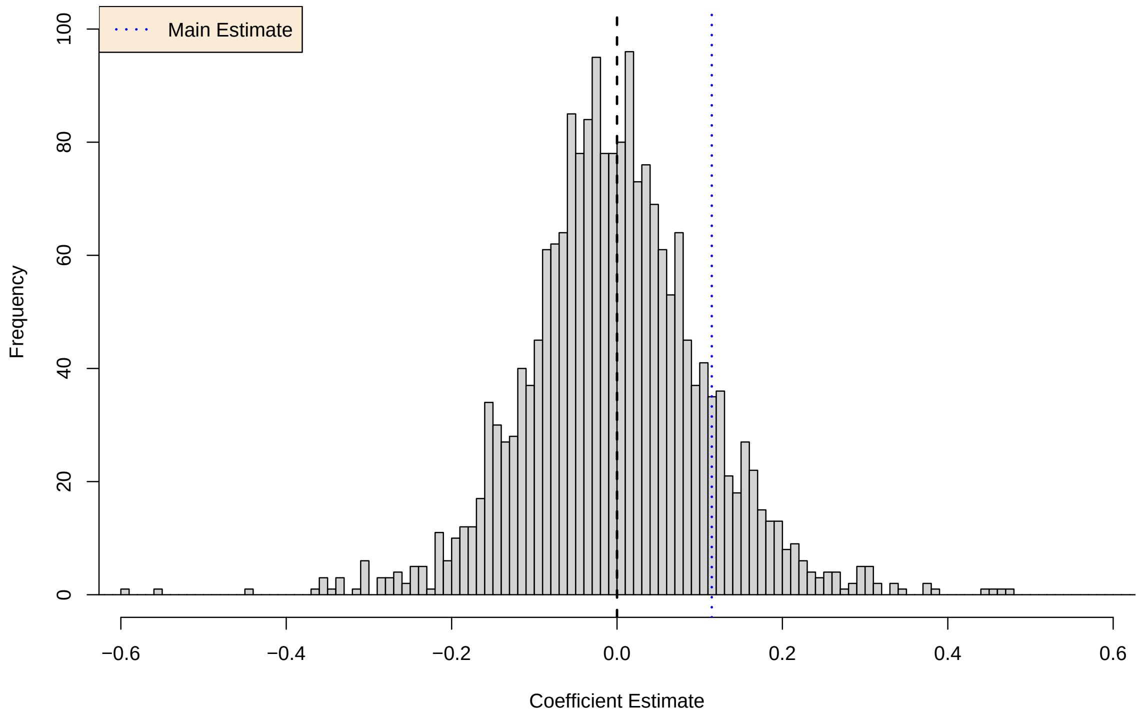

investigating whether specific brokerages react more or less to the shock of having a colleague 'winner.'

We are interested in the brokerage characteristics that are associated with an environment in which spurious shocks to individual analyst spill over to their coworkers. To do this, we first need to quantify the degree to which the shocks discussed in Section 2.3 propagate through the network of fellow analysts. We do so by estimating the following regression specification, which includes

Figure 2.9: Placebo Test – Brokerage Spillover

This figure presents coefficient estimates of our second placebo test. In this exercise, we randomly assign a proportion of winners, $\widehat{W}_{i,y}$, to brokerage-years that we know to have no winners, whilst setting the same proportion to zero for all brokerage-years with known winners. The specification is then the same as in our main analysis: $\mathbb{E}_{i,j,t}[EPS_{j,y}] = \phi_{i,j} + \tau_t + \alpha_w w_{i,y} + \alpha_l l_{i,y} + \beta_w \widehat{W}_{i,y} + \beta_l \widehat{L}_{i,y} + \Gamma X_{j,y} + u_{i,j,t}$. We repeat this 2,000 times and report the distribution of our coefficient of interest, β_w, in the figure. We find a distribution of coefficients centered around zero, with a mean of 0.00005, and a median value of 0.01159. We include our coefficient estimate from our main analysis (0.4115) as the dotted line to the right of 0. Roughly 99.65% of our placebo test coefficient estimates are lower than our main estimate.

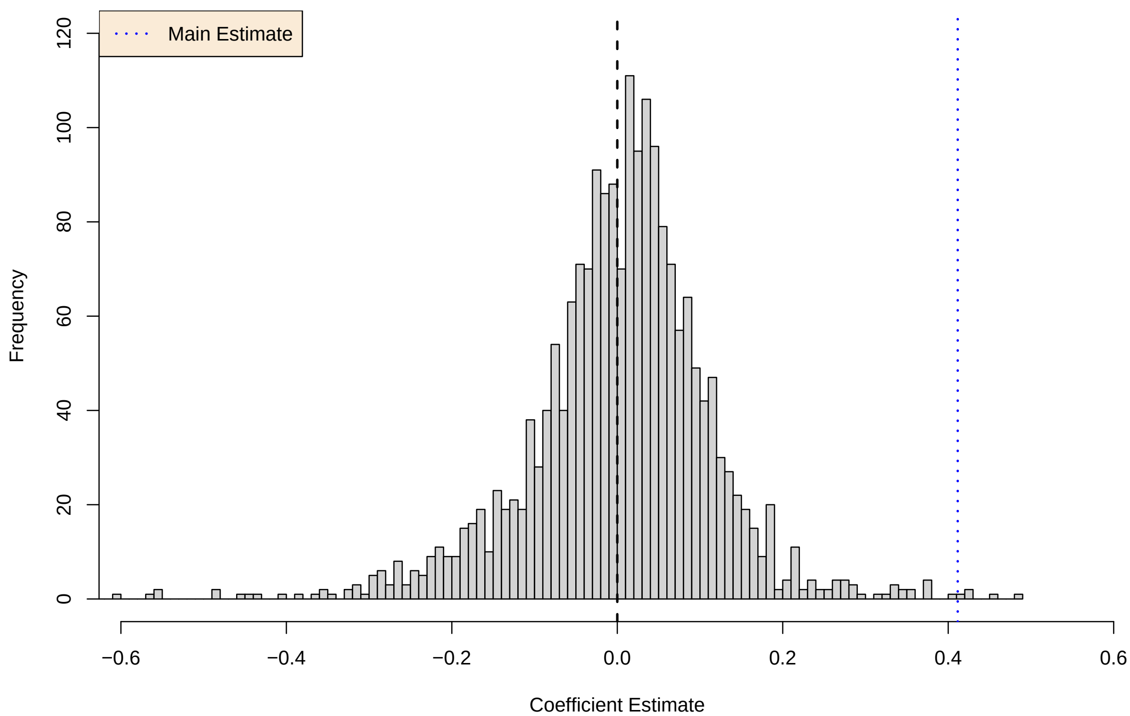

a brokerage-specific coefficient β_b:

$$\mathbb{E}_{i,j,t}[EPS_{j,y}] = \phi_{i,j} + \tau_t + \beta_0 \widehat{W}_{i,y} + \sum_{b \in B} \beta_b \widehat{W}_{i,y} \times \mathbb{1}\{i, t \in b\} + \Gamma X_{j,y} + u_{i,j,t} \tag{2.4}$$

In the specification above, i indexes analysts, j indexes firms that are covered by analysts, and b indexes brokerages.

Most importantly for us, $\mathbb{1}\{i, t \in b\}$ is a dummy variable which takes on a value of one if analyst i is working at brokerage b in year t. We restrict our attention to brokerages that have at least one non-zero observation for $\widehat{W}_{i,y}$.

Since analyst-brokerage relationships are sticky and there are a limited number of brokerages, we construct our measure (the 'Bro Score') as a time-invariant variable.[18] Furthermore, since $\widehat{\beta}_b$ is a noisy measure, we rank the coefficients and derive a relative ranking of brokerages, a measure we term 'Bro Rank.' Table 2.11 reports this ranking, the estimated 'Bro Score,' as well as the number of analysts employed at these brokerages and the unique firms covered by them.

Armed with our measures, we are interested in brokerage characteristics that are associated with particularly low/high levels of 'Bro-ness.' As discussed in the introduction, males are disproportionately more likely to describe themselves as 'avid' fans of college sports than females.[19] With this in mind, we use the gender of the analysts we were able to identify to construct the share of female analysts at the brokerages. We then test whether firms with higher 'Bro Ranks' have more or less female representation than those with lower values of the measure.

We move on to using data from MSCI/KLD to construct ESG (Environment, Social, and Governance) scores of the brokerages we have 'Bro Scores' for. ESG scores track the number of strengths and concerns that a given business has across six dimensions of ESG concerns: (1) environmental, (2) human rights, (3) diversity, (4) governance, (5) employment relations, and (6) community engagement. The score for each dimension is simply the sum of strengths within that dimension, minus the sum of concerns. Strengths and concerns are binary, taking only a value of one or zero. Overall ESG scores are constructed by simply summing the scores of all six dimensions. Using this data, we test whether brokerages with higher 'bro scores' are more or less 'responsible' in their business practices.

In a similar vein, we also use data from (Li et al., 2021) on corporate culture measures to test whether our 'Bro Scores' are associated with any of the five measures documented in that paper. Those dimensions are: (1) Integrity, (2) Teamwork, (3) Innovation, (4) Respect, and (5)

[18] The average number of brokerages an analyst works at in our sample is 1.18, the corresponding figure for IBES is slightly greater at 1.49, partly due to brokerage mergers, which are handled differently therein.

[19] 37% of men describe themselves as 'avid' fans, vs. 14% of women. These figures were reported in a recent (January 2023) poll of 2,201 American adults, published by Morning Consult, a global decision intelligence company.

Table 2.11: Bro Score & Bro Rank

This table reports the ranking of brokerages in our sample based on a measure, which we call the 'Bro Score.' This index captures the ease with which spurious shocks to one analyst at the brokerage spill over to their colleagues. It corresponds to the estimate $\widehat{\beta}_b$, which we obtain from the following regression $\mathbb{E}_{i,j,t}[EPS_{j,y}] = \phi_{i,j} + \tau_t + \beta_0 \widehat{W}_{i,y} + \sum_{b \in B} \beta_b \widehat{W}_{i,y} \times \mathbb{1}\{i \in b\} + \Gamma X_{j,y} + u_{i,j,t}$. We further report the number of analysts working at these firms and the number of unique firms covered at any point over the 2000-2022 sample period.

Rank	Name	# Firms Covered	# Analysts	Bro Score
1	Prudential Equity Group	156	15	19.84
2	Sidoti & Co	1162	179	18.53
3	Piper	54	6	16.75
4	Stifel	1314	125	16.24
5	RBC	1484	198	15.94
6	BMO	229	25	15.64
7	Oppenheimer	1307	124	15.16
8	Goldman Sachs	1792	318	14.15
9	Deutsche Bank	7	4	14.15
10	Raymond James	826	69	14.09
11	Canaccord	29	4	14.06
12	Wells Fargo	1354	158	14.03
13	AG Edwards	7	4	13.78
14	HSBC	115	55	12.56
15	Merrill Lynch	1983	437	12.52
16	Macquarie	3	2	12.52
17	Jefferies	1791	234	12.48
18	B Riley Securities	695	61	11.22
19	Credit Suisse	614	65	10.90
20	William Blair & Co	762	73	10.55
21	Morningstar Investment	820	97	9.78
22	Banc of America	839	87	8.11
23	Barclays	1274	147	2.25
24	Lehman Brothers	41	7	2.25
25	Morgan Stanley	6	2	1.86
26	Suntrust Robinson	289	18	0.79

Quality. These scores are increasing in salience, i.e. a high score for teamwork indicates the positive presence of that quality.

To assess the relationship between these brokerage-level variables and our index, we regress the measures of interest on the index as well as the log of the index.

$$measure_{b,y} = \tau_t + bro\_index_b + \Gamma X_{b,y} + \varepsilon_{b,y} \tag{2.5}$$

We are regressing the measure on the index, rather than vice versa, to exploit variation in the former over time. We control for a year fixed effect, τ_y, and include a number of brokerage-level controls: specifically, the number of analysts at the brokerage, the number of firms the brokerage covers, the number of forecasts that brokerage reports, and the average size (in total assets), book-to-market, earnings, cash holdings, and book-value-per-share of the firms that the brokerage covers.

2.5.1 Female Representation

We first look at female representation; the coefficient estimates are reported in Table 2.12. We find that the larger the 'Bro Score,' and the higher the 'Bro Rank,' the lower the proportion of female analysts working at the brokerage. The size of the coefficient is economically significant: moving from the brokerage with the lowest rank on the 'Bro Score' to the brokerage with the highest is associated with a fall in the overall proportion of female analysts of 6.24%. Given that the average proportion of female analysts within brokerages across our sample period is only 13.7%, this is a sizeable gap across the index.

We illustrate our findings further in Figure 2.10, were we plot the 'Bro Rank' and the percentage of female analysts of the brokerage along with a line of best fit. A clear upward slope is apparent, again consistent with brokerages with a higher 'Bro Rank' having fewer female analysts working at that brokerage.

2.5.2 ESG Scores

We now turn to the ESG profiles of the brokerages, investigating whether our 'Bro Score' correlates with any one of the six dimensions mentioned above, as well as with the aggregated ESG score. Our results for ESG can be found in Table 2.13. We find that brokerages with higher 'Bro Scores' have lower MSCI/KLD scores in the environmental, employment relations, and community engagement dimensions, higher MSCI/KLD scores in the human rights dimension, and lower ESG scores overall.

Table 2.12: Female Representation and 'Bro Score'

This table presents coefficient estimates from the regression of female representation as a percentage at a given brokerage on our 'Bro Rank' and 'Bro Score.' These measures capture the degree to which our college football shock diffuses between analysts working at the same brokerage. We control for the average book-to-market, size, earnings, cash and short-term asset holdings, and dividends-per-share of the firms the brokerage covers, and the number of firms, analysts, and forecasts that the brokerages make. Standard errors are clustered at the year-level.

Dependent Variable:	Proportion of Females	
Model:	(1)	(2)
Variables		
'Bro Rank'	0.0024***	
	(0.0004)	
log('Bro Score')		-0.0287***
		(0.0045)
Controls	Yes	Yes
Fixed-effects		
Year	Yes	Yes
Fit statistics		
Observations	415	352
R^2	0.28471	0.31468

Clustered (Year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

2.5.3 Measures of Corporate Culture

Finally we consider the corporate culture scores constructed in Li et al. (2021). We take the log of these scores so that our coefficients have a natural interpretation. Our results can be found in Table 2.14. Consistent with the idea of sentiment diffusion occurring amongst peers who interact day-to-day, brokerages with higher 'Bro Scores' have significantly higher teamwork scores. Perhaps surprisingly, given stereotypes surrounding sports culture, we also find that higher 'Bro Scores' correlate with higher 'Respect' scores. However, it is worth noting that this result is broadly consistent with our finding on the 'Human' component of ESG scores, as documented in Section 2.5.2.

Taken collectively, these results suggest that observable features of a workplace social network

Figure 2.10: Female Representation and ‘Bro Rank’

This figure plots brokerage ‘Bro Rank’ alongside the proportion of female analysts who work at the brokerage. We describe the construction of our ‘Bro Rank’ measure in Section 2.5. This measure is designed to capture the strength of the diffusion of the shock of winning the NCAA College Football National championship game to colleagues of analysts who were ‘winners;’ i.e. attended the winning college.

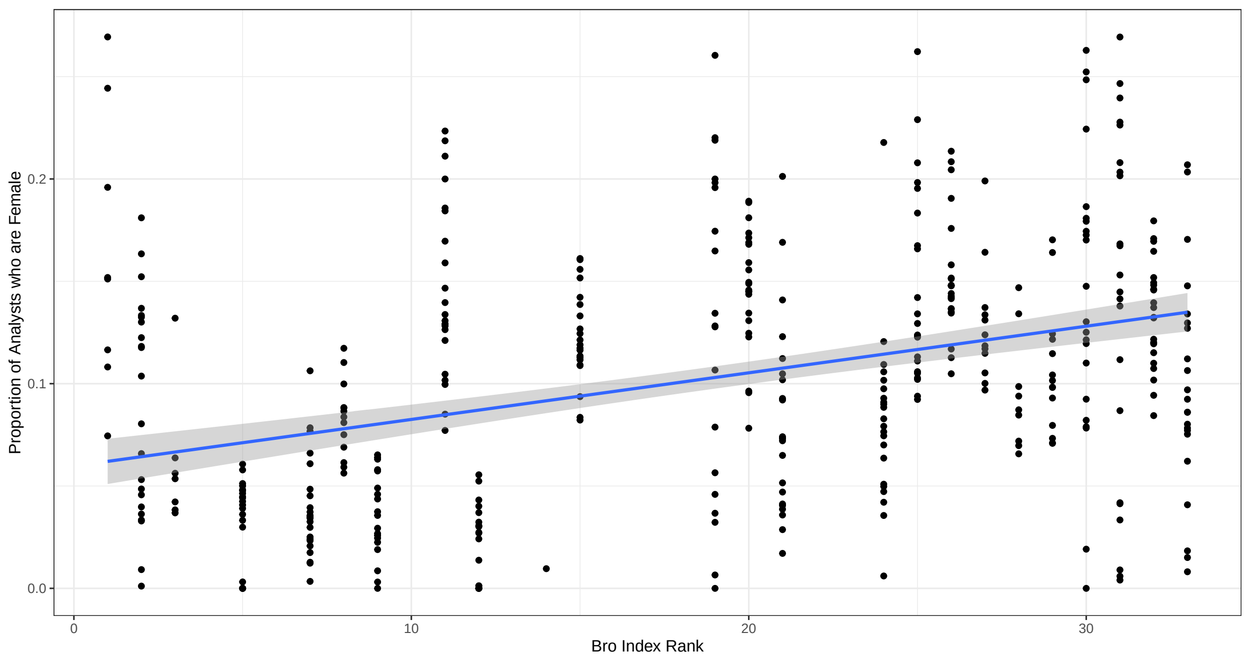

help to explain the diffusion of sentiment shocks within that network.

2.6 Firm and Market Responses

In this section, we assess whether forecast changes induced by the college football shock lead to firm or stock market level responses. We begin by investigating firm-level responses. Several papers have documented that managers engage in earnings management to meet, or attempt to meet, analyst forecasts (Almeida et al., 2016; Bhojraj et al., 2009; Shore, 2023). Using our novel source of plausibly exogenous variation in analyst forecasts, we run an instrumental variable regression to see if earnings respond in a fashion consistent with previous work. We confirm that this is the case: earnings respond close to one-to-one to forecast shocks. Subsequently, we assess how the stock market reacts to the forecast shock. Previous findings are inconclusive on the degree

Table 2.13: ESG Scores and 'Bro Score'

This table presents coefficient estimates from the regression of ESG Scores as collected by MSCI/KLD on our 'Bro Score' variable. This 'Bro Score' measure captures the degree to which our college football shock diffuses between analysts working at the same brokerage. The MSCI 'ESG' scores measure the number of strengths across various ESG dimensions net of the number of concerns in those dimensions. These ESG scores are constructed according to the six dimensions outlined by MSCI/KLD: (1) environmental, (2) human rights, (3) diversity, (4) governance, (5) employment relations, and (6) community engagement. Standard errors are clustered at the year-level. We use the same controls as in Table 2.12.

Dependent Variables:	ENV	HUM	DIV	GOV	EMP	COM	ESG
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Variables							
log('Bro Score')	-0.3068***	0.2254***	-0.1035	-0.0293	-0.2884***	-0.4083***	-1.136***
	(0.1064)	(0.0727)	(0.1484)	(0.0842)	(0.0851)	(0.0557)	(0.2854)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fixed-effects							
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics							
Observations	188	158	188	188	188	188	188
R^2	0.37386	0.33001	0.43420	0.55275	0.55842	0.42385	0.40704

Clustered (Year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

to which the stock market parses out information from analyst forecasts (Gleason and Lee, 2003; Shore, 2023; So, 2013). Consistent with Shore (2023), we find that the stock market does not react to the forecast shock. We interpret this evidence as supporting the claim that investors are able to parse out variation in forecasts driven by sentiments rather than fundamentals.

2.6.1 Firm Level Response

To identify a firm-level reaction to forecast shocks, we implement an instrumental variable design. Our identifying assumption is that the number of analysts that 'won' the NCAA National Championship game that cover a firm in a given year is orthogonal to the business conditions of

Table 2.14: Corporate Culture Scores from Li et al. (2021)

This table presents coefficient estimates from the regression of Corporate Culture Scores, taken from Li et al. (2021), on our 'Bro Score' variable. This 'Bro Score' measure captures the degree to which our college football shock diffuses between analysts working at the same brokerage. The Corporate Culture scores pertain to five key aspects of culture: (1) Integrity, (2) Teamwork, (3) Innovation, (4) Respect, and (5) Quality. Standard errors are clustered at the year-level. We use the same controls as in Table 2.12.

Dependent Variables: Model:	log(Integrity) (1)	log(Teamwork) (2)	log(Innovation) (3)	log(Respect) (4)	log(Quality) (5)
Variables					
log('Bro Score')	0.1823	0.3977^{***}	0.2326^{***}	1.304^{***}	-0.1706
	(0.1296)	(0.1335)	(0.0787)	(0.1075)	(0.1059)
Controls	Yes	Yes	Yes	Yes	Yes
Fixed-effects					
Year	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	375	375	375	375	374
R^2	0.18676	0.17147	0.35269	0.32905	0.13029

Clustered (Year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

that firm. As such, the instrument that we use is precisely that number, which we label $W_{j,y}$.[20] Let $A_{j,y}$ be the set of analysts that cover firm j in year y. Then $W_{j,y}$ is defined in the following way:

$$W_{j,y} = \sum_{i \in A_{i,j}} w_{i,y} \tag{2.6}$$

To test for the relevance of our instrument, we begin by conducting a first stage regression of our instrument on the consensus earnings forecast for a given firm-year, controlling for a firm and a year fixed effect, as well as several firm-year level covariates. Here we make use of the IBES Summary dataset to collect the IBES consensus forecasts; this is the consensus forecast that is typically

[20]We choose the sum of winners rather than the proportion in this case, as our matched dataset of analysts to college attendance is only a subsample of the universe of analysts. Given that we are interested in looking at reactions to changes in the consensus forecast, which naturally includes the universe of non-stale analyst forecasts, we view the sum of 'winners' covering a firm to be a more appropriate measure of firm exposure, as it does not rely as explicitly on the size of the matched sample. In a robustness check, we run the same exercise using the proportion of 'winners,' and find similar point estimates, albeit with a weaker first stage, and hence limited statistical significance. This is not all together surprising, given the aforementioned concerns surrounding our use of a subset of forecast data.

used for market tests (Bartov et al., 2002; Brown, 2001; Lim, 2001). We use the most recent consensus forecast prior to the forecast period end-date as the measure that earnings performance are compared against; again, this is precisely because market tests are typically performed relative to this measure. We use the mean forecast (IBES Summary item 'MEANEST') in our main analysis, although the results are near identical if we instead use the median forecast (IBES Summary item 'MEDEST'). As in the exercises conducted above, we standardize the consensus earnings forecast, and the earnings variable itself, this time at the firm level, to avoid problems associated with scale. We then run a standard instrumental variable regression.

Our findings are presented in Table 2.15. We report results for both unadjusted values of the consensus forecast and firm-level earnings-per-share, as well as for the standardized variables. In both cases, we find an F-statistic that is above the Stock and Yogo (2002) cutoff of 15, meeting the requirement of strong instruments. We find that for a one standard deviation increase in the consensus forecast, earnings increase by 0.8019 standard deviations, and we cannot rule out a one-to-one relationship.

2.6.2 Asset Market Response

While there is ample evidence that the information contained in analyst forecasts can move markets (Gleason and Lee, 2003; So, 2013), in a fully-rational benchmark without informational frictions, investors would be able to filter out any noise in these forecasts and form beliefs based exclusively on the relevant new information. If this was the case, any spurious increases in forecasts should leave asset prices unaffected.

To address the question of how changes in forecasts that are unrelated to company fundamentals affect markets, we conduct another instrumental variable exercise. Our approach is broadly in line with the procedure detailed in Section 2.6.1, though we substitute firm returns for firm earnings as our dependent variable. Again, our identifying assumption is that the number of analysts that won the NCAA National Championship game that cover a firm in a given year is orthogonal to the business conditions of that firm.

Table 2.15: IV Estimation – Firm Response

This table presents our instrumental variable regression of firm-level earnings, $EPS_{j,y}$, on the consensus analyst earnings forecast, $\mathbb{E}_{j,y-1}[EPS_{j,y}]$, where the consensus forecast is instrumented using the number of analysts who cover firm j in year y who attended the college that won the NCAA National Championship Football game in the year y, $W_{j,y}$. We control for a firm and a year fixed effect, market value, book-to-market, lagged assets, stock price, and dividends per-share, and we cluster standard errors at the firm level. The consensus forecast is taken from the IBES Summary dataset.

Dependent Variables:	$\mathbb{E}_{j,y-1}[EPS_{j,y}]$	$EPS_{j,y}$	Standardized $\mathbb{E}_{j,y-1}[EPS_{j,y}]$	Standardized $EPS_{j,y}$
IV stages	First	Second	First	Second
Model:	(1)	(2)	(3)	(4)
Variables				
$W_{j,y}$	0.0648***		0.0235***	
	(0.0124)		(0.0060)	
$\mathbb{E}_{j,y-1}[EPS_{j,y}]$		0.9467***		
		(0.2513)		
Standardized $\mathbb{E}_{j,y-1}[EPS_{j,y}]$				0.8019***
				(0.2422)
Controls	Yes	Yes	Yes	Yes
Fixed-effects				
Firm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	52,053	52,053	51,710	51,710
R^2	0.70372	0.54320	0.98575	0.97609
F-test (1st stage)	57.443	57.443	17.216	17.216
Wu-Hausman, p-value		0.55886		0.38183

Clustered (Ticker) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Following Carhart (1997) and Fama and French (2015) we compute the exposure of all stocks to the four most common risk factors. Using a rolling monthly estimation procedure, we estimate factor loadings by regressing the excess return on stock j at time t on the following factors:

$$r_{j,t} - r_{rf,t} = \alpha_j + \beta_{j,m}(r_{m,t} - r_{rf,t}) + \beta_{j,smb} SMB_t + \beta_{j,hml} HML_t + \beta_{j,mom} MOM_t \tag{2.7}$$

The rolling window that we select is the 60 months prior to the month of the return. Using the coefficient estimates $\left(\widehat{\beta}_m, \widehat{\beta}_{smb}, \widehat{\beta}_{hml}, \widehat{\beta}_{mom}\right)^\top$ we compute the abnormal return of stock j by subtracting the expected from the realized return.

$$AR_{j,t} = r_{j,t} - \mathbb{E}[r_{j,t}] \tag{2.8}$$

We then run our instrumental variable regression using both raw returns, as well as our constructed abnormal returns, as the dependent variable. We control for a firm fixed effect, an industry-month-year fixed effect, and a host of firm-year level covariates. See Table 2.16 for the coefficient estimates.

Our results are consistent with the stock market correctly interpreting the forecast shock as spurious; we fail to find evidence that the shock to the forecast moves either raw or abnormal returns, despite high power in our first stage (our first stage F-statistic is 194.96). This finding is consistent with results in Shore (2023) who also shows that the stock market does not react to plausibly exogenous variation in analyst forecasts.

2.7 Discussion

The relevance of social networks has received increased attention in the Economics and Finance literature recently. However, analyses have been limited to the spillover of potentially valuable and actionable information through networks, constituting a form of peer learning. To the best of our knowledge, we are the first to provide evidence of spurious information cascading through a network of professionals, which in turn act on said sentiment in a high stakes environment.

Bailey et al. (2018a,b) are among the first to provide tangible evidence on the relevance of social ties, which they measure using the social connectedness between individuals across regions derived from the Facebook social graph. Specifically, they find that house price expectations are very much driven by social interactions. These expectations then influence both the choice between renting and owning, as well as leverage when obtaining a mortgage. While it can be argued that

Table 2.16: IV Estimation – Asset Market Response

This table presents the results of estimating an instrumental variable regression of monthly asset returns on the standardized consensus analyst earnings forecast ($\mathbb{E}_{j,y-1}[EPS_{j,y}]$), where we standardize at the firm-level. We show results for two measures of asset returns: raw returns, $r_{j,t}$, of firm j in month-year t; and abnormal returns, $AR_{j,t}$, of firm j in month-year t, where abnormal returns are constructed based on a Carhart (1997) four factor model. We outline our procedure for constructing these abnormal returns in Section 2.6.2. We instrument for changes in the consensus forecast using the variable $W_{j,y}$, which measures the number of analysts that cover firm j in the year y who attended the college that won the NCAA College Football Championship Game in the year y. We control for estimated rolling betas, $\{\beta_0, \beta_{mkt}, \beta_{smb}, \beta_{hml}, \beta_{mom}\}$, book-to-market, price, assets, cash and short-term asset holdings, dividends-per-share, and past volatility.

Dependent Variables:	Raw Returns	Abnormal Returns	Raw Returns	Abnormal Returns
Model:	(1)	(2)	(3)	(4)
Variables				
$\mathbb{E}_{j,y-1}[EPS_{j,y}]$	-0.0015	-0.0046		
	(0.0066)	(0.0061)		
Standardized $\mathbb{E}_{j,y-1}[EPS_{j,y}]$			-0.0027	-0.0083
			(0.0120)	(0.0112)
Controls	Yes	Yes	Yes	Yes
Fixed-effects				
Industry-Month-Year	Yes	Yes	Yes	Yes
Firm	Yes	Yes	Yes	Yes
Fit statistics				
Observations	499,768	499,768	497,898	497,898
R^2	0.44844	0.28842	0.44959	0.28827
Within R^2	0.02953	0.01314	0.02939	0.01178
F-test (1st stage)	290.37	290.37	194.96	194.96
Wu-Hausman, p-value	0.46141	0.25112	0.52634	0.28797

Clustered (Industry-Month-Year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

the experiences of (geographically) distant friends are not directly applicable to the local housing market, there unarguably is portable knowledge that is applicable across different national housing markets. Their findings therefore are different from ours in two important ways. Firstly, while real estate is the most important asset in the portfolio of most households, making the purchase decision a very high stakes one, households can still be considered unsophisticated market participants potentially putting too much weight on information obtained through their social network. In contrast, the individuals in our setting are professionals and thus, arguably, less susceptible to irrelevant information. Secondly, as argued above, if there are indeed common factors across housing markets, relying on the experiences of others, even in distant locations, can improve the accuracy of expectations. In our setting, on the other hand, the shocks are orthogonal to information relevant to the decision at hand, making it all the more puzzling that individuals react to them in the first place.

Further evidence on the reaction of professionals to outside information has been provided by Kempf and Tsoutsoura (2021). In considering credit analysts, their setting is similar to ours in that the individuals under consideration have extensive experience in financial markets and have to make decisions that influence assets worth many millions to billions of dollars. However, in contrast to our sentiment shock, the authors consider the impact of partisanship on the analysts' decisions. Specifically, they compare credit rating made by analysts that are registered with either the democratic or republican party around changes in presidency. While their results are similar to ours, i.e. democrat (republican) analysts give better ratings when a democratic (republican) candidate takes office, it is hard to argue that the color of the white house is unrelated to the corporations they cover. There obviously is reason to believe that these reactions are driven by the different expectations regarding the future wellbeing of the economy. It thus cannot be concluded that any kind of reaction is spurious. Our setting is different from the aforementioned ones in that the shock we consider is orthogonal to firm fundamentals and entirely sentiment-driven in nature.

Similarly, our results are different from the recent literature on the reaction of sell-side analysts to local geographic shocks. For example, Cuculiza et al. (2021) focus on the impact of terrorist attacks on earnings forecasts, but do so at a local level; they find that analysts who are closer to a terrorist event are affected more strongly than analysts located farther away. Similarly, Kong et al. (2021) show that proximity to earthquakes lowers analysts' optimism. Whilst consistent with our findings that analyst forecasts are influenced by personal shocks that are plausibly independent of the business conditions of the firms they cover, it is difficult to see how these broad shocks could be used to identify network spillovers, precisely because the shock affects the individual as well as her network.

2.8 Conclusion

In this paper we provide evidence that analysts not only react to shocks that are specific to themselves, but also show that such shocks spill over to other analysts, specifically their colleagues, through their social networks. We do so using a novel, hand-collected dataset on the college attendance of 7,481 analysts over the 2000-2021 period. To the best of our knowledge, the shock we use, whether an analyst's college football team wins the NCAA finals, similarly is new to the literature.

Our findings support the claim that analyst forecasts are subject to bias and sentiment in much the same fashion as many economic decisions. This is surprising since the analysts in our sample are professionals engaged in a high stake setting, with many billions of dollars of assets being managed following their guidance. In this paper, we also show that analysts are influenced by the sentiments of agents in their social network, consistent with evidence from related work. While peers naturally react less strongly to these shocks, the effects are statistically significant and economically meaningful. Being able to trace out the flow of information between equity analysts is a major novelty in our approach.

With respect to these network effects, we find that the ease with which these spurious shocks propagate to other analysts at the same brokerage correlates with observable characteristics of the work environment. Brokerages where the degree of this diffusion is greater have lower female representation in their analyst teams, as well as lower ESG scores. Finally, we substantiate existing evidence that analyst forecasts causally influence the economic decision making of firms, while only finding limited support for a similar relationship with investors.

Chapter 3:

Sharing is Caring?

Knowledge Diffusion in Researcher Networks[1]

3.1 Introduction

Innovation is widely regarded as one of the key drivers of economic growth and thus has received significant attention from researchers and policy makers alike.[2] As the world has become increasingly complex and technology more advanced, most groundbreaking innovations are no longer attributable to innovators working in solitude (such as Edison), but ever-growing research teams. Figure 3.1 depicts the evolution of the share of U.S. patents developed by teams of inventors. More than three quarters of patents granted in the U.S. in 2021 were jointly filed by multiple inventors. Developing an understanding of how interactions between researchers shape the research and development (R&D) process is therefore crucial considering these recent developments.

Peer effects and social networks have been identified as playing a key role in many areas of economics. Yet, their relevance in research and development has not been explored in detail, due to the lack of sufficiently granular data. Private sector R&D is subject to a lot of secrecy, making it nearly impossible to observe any form of interactions between inventors. As a substitute, research in innovation economics usually relies on patent counts and citations, measures that are naturally biased towards successful research projects (Hall et al., 2001). Only recently have the identities of inventors receiving increasing attention (Kim and Marschke, 2010), however, by still relying

[1]This chapter is based on Fischer (2022). I am grateful for valuable discussions with Tania Babina, Xavier Giroud, and Paul Tetlock. This paper has furthermore benefited from comments by seminar participants at Columbia Business School.

[2]The United States started granting a tax credit for qualifying research expenditures in 1981 with the goal of boosting investment in R&D. Most recently, the CHIPS and Science Act of 2022 will provide $280 billion in additional funding for domestic research in semiconductors.

Figure 3.1: Inventor Teams

This figure depicts the share of patents granted in a given year, which have been developed by more than one inventor. The data is restricted to patents filed in the United States and has been constructed from the USPTO Patent Database. The share of multi-inventor patents has steadily increased since 1974, growing from 40% to 75%.

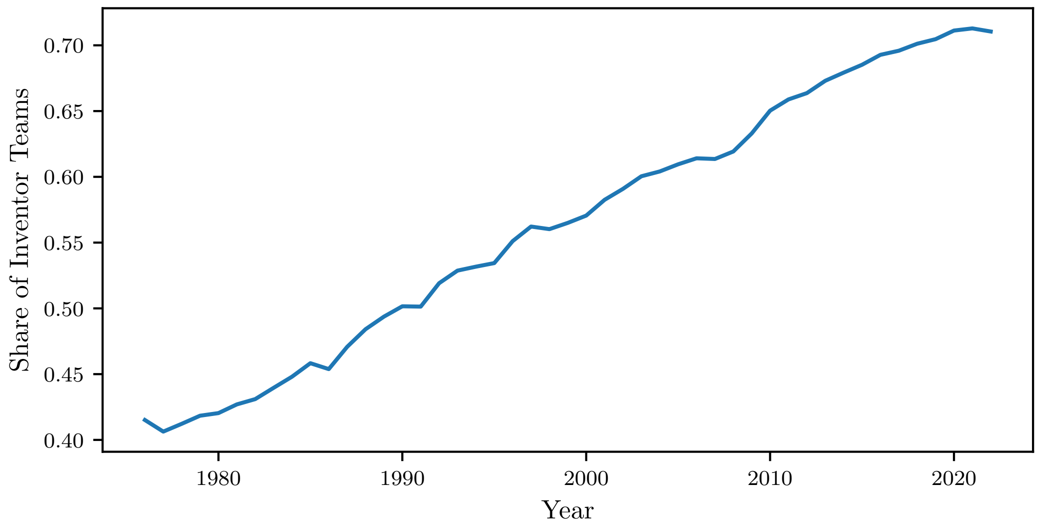

on data from the U.S. Patent and Trademark Office (USPTO) any such research abstracts away from interactions between inventors, and solely focuses on research outcomes.[3] Due to these shortcomings, novel datasets are starting to grow in relevance in uncovering the nature of human interactions in the research process. At the forefront of this development are Atkin et al. (2022), who use anonymized phone tracking data to proxy for interactions across firms in Silicon Valley, relating a proxy for casual meetings to patent outcomes.

Instead of relying on anonymized data from the private sector, this paper uses data on interactions between academics on Twitter, relating measures of connectedness to research outcomes. The use of data from social media has become increasingly common in studying various kinds of interactions between economic agents at the micro-level. Most prominently, Bailey et al. (2018a,b, 2022) use data from the Facebook social graph to shed light on peer effects in the housing, lending, and product markets. More closely related to the approach taken in this paper, Cookson and Niess-

[3]This still holds when relying on patent filings instead of data on patents granted, as the filing itself will only be observable for successful research projects.

ner (2020) and Cookson et al. (2023) document the nature of disagreement in financial markets and presence of echo chambers among professional investors on a Twitter-like social network.[4]

In considering academics as stand-ins for researchers more broadly, and private-sector inventors specifically, this paper follows the spirit of Azoulay et al. (2011).[5] While academic research is distinct from R&D in the private sector, it is significantly more transparent while still featuring a high degree of quantifiability of research outcomes, addressing some of the concerns raised in the context of the existing literature. Linking researchers to their Twitter accounts, interactions between individuals become observable and can be related to measures of productivity and research outcomes. The increased importance of Twitter in facilitating communications between economists (#EconTwitter) during the COVID-19 pandemic gives rise to a laboratory in which the role of interactions in the research process can be analyzed. Using the network structure inherent to Twitter allows for the development of a much better understanding of how knowledge diffuses through the network of economists and how it impacts future research.

Using data from the Social Sciences Research Network (SSRN), Research Papers in Economics (RePEc), and the National Bureau of Economic Research (NBER), a sample of 1,072,310 research projects co-authored by 145,481 unique economists is constructed. Of those individuals 19,422 are linked to their Twitter accounts. Based on more than fifty million unique tweets, interactions between individuals are quantified. Tweet text and metadata are used to assess the types of interactions between individuals. Specifically, tweets relating to academic research can be linked to the research projects mentioned therein.

The findings in this paper are both qualitative as well as quantitative in nature. On the qualitative side, a strong presence of homophily in the discussion of research on Twitter is documented. While the degree to which preference is given to in-group sources differs between research areas, it is pronounced across the board. In the sample, social interactions center around, in order of importance, news, ongoing research (i.e., working papers), and completed research (i.e., journal articles).

[4]Stocktwits is a social network that allows users to share their views on different financial instruments, predominantly equities.

[5]Further papers shedding light on the research process by considering academia as a laboratory are Agrawal and Goldfarb (2008) and Murray et al. (2016) among (many) others.

This ordering speaks to the relative importance of interactions at different points in the research process. As projects mature, the kind of interaction changes from gathering insights to marketing the research project. At this stage of the research process, individuals self-promote their own work as well as endorse that of others. While there are some that engage in both, most individuals cluster at the extremes. The relative importance of these motives differs between fields.

Informed by these observations, proxies for interactions among individuals are related to research outcomes (productivity, visibility, and impact) in more quantitative, reduced-form analyses. An increase in social interactions is associated with greater productivity and visibility, but a reduced impact, as measured by forward citations. However, as the effect is contemporaneous, it is difficult to rule out reverse-causality. Using time-stamped download data in an event study, the causal effect of working papers appearing on the social network on the number of downloads they receive is assessed. At an associated increase in lifetime downloads of more than 20%, the effect is economically meaningful. This finding suggests the importance of visibility of research, especially early on in the roadshow. Splitting the sample into cases of (self-) promotion and endorsement, it is found that this effect is stronger when individuals other than the authors reference a working paper. This is consistent with readers being sophisticated and controlling for the authors incentives when sharing their work. When shared by a third party, these concerns are less relevant, leading to a stronger effect. These findings furthermore serve as suggestive evidence of a trade-off between improving research projects (and increasing their impact) and marketing them (increasing visibility) when the authors' time is constrained.

Related Literature. This paper contributes to several stands of literature. Firstly, it adds to the large body of work on research and innovation. As a key driver of economic growth, innovation and its production have received a great amount of attention (Hall, 1995), most of which is empirical and relies on patent data to measure outcomes (Hall et al., 2001). Naturally, the evaluation of R&D incentives is one of the key questions that has been asked (Dechezleprêtre et al., 2023; Hall and Reenen, 2000; Wilson, 2009). Another branch of the literature explores the determinants of becoming an inventor (Aghion et al., 2017; Bell et al., 2019) and the relevance of teams in research (Kim and Marschke, 2010). Focusing on competition, Bloom et al. (2013) and Arora et al. (2021) explore incentives to innovate in competitive markets. Most existing papers rely on balance sheet or patent measures to quantify the investment in and outcomes of research. In doing so, they abstract away from the actual innovation process and consider output measures that are likely biased, since only successful interactions between inventors are observable. To alleviate such concerns, novel datasets are key to understanding how interactions shape innovation production. Most closely related to this paper is the work by Atkin et al. (2022) who use anonymized geolocation data to identify interactions between employees from different firms in Silicon Valley to speak to the relevance of these meetings for patent citations and other spillovers.

In addition to the large empirical literature, there are various papers that propose models for different aspects of the research process. Relevant to the setting in this paper are those that specifically consider priority races as well as the interactions between individuals. With respect to the former, Bobtcheff et al. (2016) model the competition between two researchers for a prize that will be awarded to the winner of the race. Hill and Stein (2021) generalize their model and provide evidence from the natural sciences. The underlying economics explored in this paper are most closely aligned with the model proposed by Stein (2008), in which cumulative innovation occurs if competing inventors interact.

Furthermore, this paper is related to the large literature on information diffusion and peer effects in social networks. The relevance of peer learning has been established in a wide range of settings. Cohen and Frazzini (2008) and Cohen et al. (2008) consider the flow of information in financial markets, specifically through education networks between corporate insiders and analysts/fund managers. Similarly, Shue (2013) uses data on the random assignment of MBA students to courses to address the relevance of social networks in business decision making. Similarly, Fischer and Shore (2023) consider the spread of non-financial sentiment shock within brokerages. Recently, data from social media platforms has seen an increased use in answering such questions. Among the first and most prominent examples thereof is the empirical work by Bailey et al.

(2018a,b, 2022). Using the Facebook social graph, the authors relate connections between geographic regions in the U.S. to various outcome measures (house prices, mortgage leverage choice, product adoption), identifying peer effects as an important channel. Using a dataset like the one in this paper, Cookson and Niessner (2020) and Cookson et al. (2023) document the nature of disagreement in financial markets and presence of echo chambers among professional investors. Why economic agents would want to share profitable ideas is the subject of the analyses performed by Crawford et al. (2017).

Lastly, by considering academia as a laboratory for the research and innovation process, this paper naturally relates to the body of work on academic research. Azoulay et al. (2010) consider the impact of the unexpected death of superstar scholars on their coauthor network using data from the natural sciences. With respect to interactions between academics, Agrawal and Goldfarb (2008) use the introduction of ARPANET, a precursor of the internet, at universities to assess the impact of a reduction in communication costs on co-authorship. In addition to serving as stand-in for inventors in the private sector, the process of academic research itself has been of great interest. Aghion et al. (2008) model the trade-off between academic freedom and private sector focus in shaping the process of innovation. More recently, Babina et al. (2020) focus on the origin of founding sources and how they affect research outcomes. With respect to academic research in Economics, Brogaard et al. (2014, 2018) consider the relevance of proximity to editors and the tenure system, respectively.

The remainder of this paper is organized as follows: Section 3.2 describes the datasets used, the variable construction, and presents summary statistics. The empirical approach is outlined in Section 3.3. Qualitative and quantitative results are presented in Section 3.4 and Section 3.5 respectively. Ultimately, Section 3.6 concludes.

3.2 Data

To assess the relationship between social interactions and research outcomes a novel dataset is constructed. First, the setting which will be used as a laboratory is outlined. Subsequently, the exact data sources and the process used to combine them is described. Finally, summary statistics for the different samples used in this paper are presented.

3.2.1 Setting

The R&D process in the private sector is intentionally secretive and opaque to ensure that competing inventors/teams/firms do not obtain access to ideas before they can be protected by a legal right (Arora et al., 2021). As such, answering questions related to the identities of and interactions between inventors is difficult. In an ideal experiment, one would record all interactions between inventors with their peers, gathering information on the types of interactions they have. Since this is not feasible, two abstractions are made in an attempt at creating a meaningful laboratory in which the relationship between interactions and research outcomes can be studied. Abstracting away from the interactions themselves and considering their outcomes instead, previous studies have focused on patent inventors (Kim et al., 2006; Kim and Marschke, 2010). The selection of individuals into research roles is analyzed by Bell et al. (2019). More recently, Atkin et al. (2022) overcome the aforementioned issues by using anonymized geolocation data of employees in Silicon Valley to identify potential meetings across firms and link them to innovation measures. While a significant improvement over the prior literature, the 'meetings' identified in their study can only serve as a proxy for general interactions between firms and not those between R&D teams specifically. There currently is no work on the actual context of interactions between such researchers and how they relate to research outcomes, a gap this paper seeks to fill.

Rather than focusing on inventors in the private sector, academics in Economics are used as a laboratory to evaluate different theories regarding the relevance of human interaction in the research process. This motivation is in line with Azoulay et al. (2011), who argue that academic research is a tractable form of innovation that is more readily observable and less prone to secrecy. Furthermore, Economics as a quantitative and quantifiable field, which is reasonably isolated from other areas, reducing spillover effects, is a natural starting point.[6] The research process is surprisingly similar to that in the private sector. Unlike many other fields, the ultimate goal for every paper is to be published in a renowned journal, a process not dissimilar from applying for a patent. Furthermore, in contrast to many other fields, pre-prints (i.e., working papers) and presentations at seminars/conferences play a key role, alleviating concerns that research projects only become observable upon completion.

Ideally, one would like to tape conversations between individuals (e.g., at major conferences) to obtain data on interactions and their content. For obvious reasons, this is not feasible, necessitating a second layer of abstraction, the use of social media data. Over the past few years, Twitter gained popularity in communicating economics research, a phenomenon commonly referred to as #EconTwitter. Despite playing a key role in Economics today, the term itself lacks a concise definition. For the purposes of this paper, it is defined as any activity by research economists on Twitter that, broadly speaking, pertains to topics in Economics. Figure 3.2 provides a visual representation of this definition, which will serve as guiding principle in the data collection. Nevertheless, three major concerns remain.

Firstly, one must question whether the interactions between economists on Twitter can proxy for conversations that would also happen in different settings. During the COVID-19 pandemic, when in-person interactions such as major conferences like the annual meeting of the American Economic Association (AEA) were conducted online, activity on Twitter increased and interactions shifted from a discussion of news to more directly research-related areas. Anecdotal evidence suggests that the interactions on #EconTwitter resembled those otherwise only observed in conference-like settings (see acknowledgements in Anton et al., 2021).

Secondly, secrecy will always remain a concern when it comes to the discussion of in-progress

[6]Fourcade et al. (2015) document the isolatedness of economics, with less than 5% of citations in the top five economics journals being made to other fields. The only relevant related field is Finance, which is going to be grouped into Economics in this paper.

Figure 3.2: What is #EconTwitter?

For the purposes of this paper, #EconTwitter is defined as any activity by research economists on Twitter that pertains to topics in Economics. The sample of economists is constructed based on authors of working papers uploaded to the SSRN *Economics Research Network* (ERN) and *Financial Economics Network* (FEN), as well as individuals with profiles on *Research Papers in Economics* (RePEc) or the *National Bureau of Economic Research* (NBER). Authors are subsequently linked to their Twitter profiles. Whether Twitter activity is related to topics in economics is determined based on the links embedded in tweets.

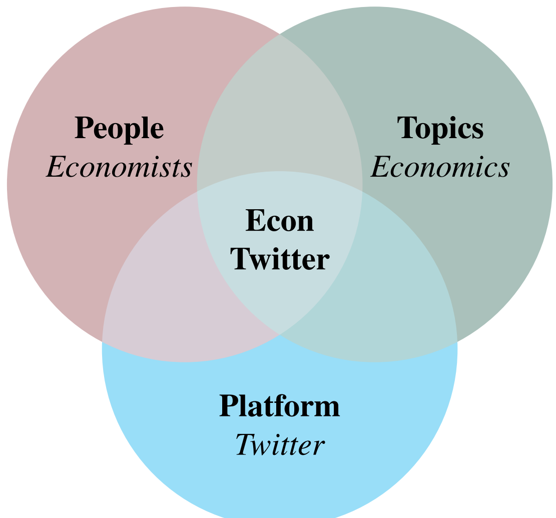

research, be it in the private sector or in academia. For the aforementioned reasons, this problem is expected to be less pronounced in academic research. With working papers playing a significant role in research in Economics, as well as comments on such papers being actively encouraged by the profession, part of the overarching concern can be alleviated.

Lastly, it is an open question whether economists on Twitter are representative of the population of scholars in the field. Of the 1,625 members of the National Bureau of Economic Research (NBER), 716 have active Twitter accounts. While this subsample is of course biased towards more influential economists, these figures do speak to the large presence of economists on Twitter. As will be demonstrated in later sections, when controlling for tenure/skill, individuals with and without Twitter accounts are remarkably similar.

3.2.2 Data Sources

As described in the previous section, #EconTwitter will serve as a laboratory to analyze the relevance of interactions on research productivity and impact. The starting point in constructing the sample is the identification of research economists and their publication records. To not bias the sample towards more prominent scholars ex-ante, a wide range of sources is used. As it is at the authors' discretion which repository to add a working paper to, data from all sources is compiled to obtain a comprehensive list of research projects.

RePEc. Research Papers in Economics (RePEc) is an open-source project seeking to facilitate the distribution of Economics research. The project directly collects publication data from journals in economics and a selection of institutions that feed their working/policy papers to the website. Each research paper is assigned a unique identifier, allowing to track the evolution of papers over time. The entirety of the RePEc database is downloaded, which contains 3,544,637 papers by 79,364 authors.

SSRN. The Social Science Research Network (SSRN) is Elsevier's preprint service which has become the major platform to share working papers in many social sciences, including Economics. All working papers that are part of either the Economics (ERN) or Financial Economics (FEN) network are retrieved using the SSRN API (see Appendix) and matched to the author's current affiliations and rankings (metadata from SSRN). Once a paper has been published, the journal it appeared in is recorded as well. Before any further filtering, this part of the sample consists of 631,683 papers by 381,076 authors. To zero-in on individuals actively engaged in academic research, all authors with only a single paper are dropped, reducing the sample to 147,882 authors of 632,107 papers. In addition to the number of downloads at the time of collecting the data, daily downloads for 37,305 working papers that are mentioned in tweets.[7] There are 5.7 million daily observations with a non-zero number of downloads, or an average of 151.23 observations per paper.

[7]This data was curated by SSRN/Elsevier for this project and is not publicly available.

NBER. As the National Bureau of Economic Research's (NBER) working paper series is among the most well-known in the profession, all entrants and their authors are obtained. Viewed in isolation, the NBER provides another 30,151 working papers written by 14,962 authors.

Twitter. The Twitter API is used to obtain the full history of tweets, retweets, and comments for the identified accounts. For each of them, the full metadata, including timestamps, is downloaded alongside the text. Furthermore, the metadata for the accounts is collected as well, which includes the user's biography, sign-up date, and anything else visible on the account page.

Miscellaneous. To quantify the visibility and impact of academic work, citation data is obtained from SSRN, Crossref (2021), and Google Scholar.[8] The SSRN data contains the total number of citations received by all SSRN working papers, as well as the number of times these have been downloaded from the website but lacks a time-series dimension. Citation data from Crossref and Google Scholar is dated, allowing for the construction of the impact of scholarly work over time.

After collection, the data is combined and disambiguated, such that every project, from the first working paper to the final publication, as well as all authors are uniquely identified. As working papers can appear repeatedly in various sources, the titles and lists of authors are used to identify records that correspond to the same research project. Similarly, many authors have multiple accounts, both across as well as within the three sources. Using names and affiliations, each author is assigned a unique identifier. These two processes, on the paper and author side, are conducted jointly, using the link between papers to validate the author-link and vice versa.

Lastly, economists are linked to their Twitter accounts. An appealing feature of the NBER and RePEc data is that a large number of researchers disclose their Twitter accounts in their profiles.

[8] Crossref is the major DOI registration agency and specializes in academic articles. Publishers as well as aggregators, such as Google Scholar, rely on the linkages provided by Crossref to count citations and enable linking to cited articles.

Using the Twitter API, these accounts are checked for validity and a link is established.[9] For individuals who list their personal website, these are scraped and checked for links to a Twitter profile with a matching username. In a final step, Twitter profiles with names similar to all unmatched academics are obtained using the Twitter academic API. Using information on affiliations, these potential matches are then automatically validated. At every step of the way, matches are identified conservatively. For each identified economist, the entire history of tweets as well as the current list of accounts they are following and are followed by are obtained.

3.2.3 Variable Construction

One key contribution of this paper is the identification of economist's Twitter accounts relative to the existing literature. Bisbee et al. (2020), who look at political scientists, start from a list of 131 U.S. universities, hand-matching all faculty to their Twitter accounts. Unlike this highly manual approach, Filippas and Horton (2021) use self-reported Twitter accounts on RePEc and identify other accounts belonging to #EconTwitter based on followership. This limits the ability to include outside data, such as citation records, as individuals are not identified/disambiguated.[10] In contrast, this paper combines data on the Twitter accounts of academics with measures on their research output for an entire field, alleviating concerns regarding sample selection.

Research Areas. Using data on the Journal of Economic Literature (JEL) classification of all papers in the RePEc sample, individual authors are assigned to one of the twenty economic subfields (by JEL codes). In assigning authors to research areas, the area most commonly associated with an individual is picked, to control for coauthors' specializations. A list of all JEL codes and the corresponding research areas is provided in Table C.1.

Social Interactions. Data from Twitter is used to measure social interactions. For each indi-

[9] As both RePEc and the NBER identify Twitter accounts using user handles, which can change over time, it is important to ensure that these are indeed existing accounts.

[10] Since Filippas and Horton (2021) are concerned with the content production on social media, they do not require any external data.

vidual, the year they joined the platform is computed ($joining\_year$). This naturally gives rise to a dummy variable that captures whether an economist had an account in a given author-year ($had\_account$). Using data from more than fifty million tweets, the variable $tweeted$ takes on a value of one if the individual tweeted at least once in a given year. In addition, the intensive margin of activity is captured by the logarithm of the total number of tweets in a year ($\ln(1 + tweets)$).

Productivity. The number of research papers written is a natural measure of productivity ($\#papers$). It is analogous to the filing of a patent, commonly used as a similar measure in the innovation literature. Since working papers are revised between their emergence and publication, it is important to be able to trace back each revision to the posting of the initial version. This implicitly carries the assumption that only new projects are considered for the purposes of productivity. In this paper, productivity is quantified as a simple count over projects an author engages in, regardless of the number of coauthors on a given paper. An alternative approach would be to weigh projects by the number of authors. Since the scope of research projects varies widely across as well as within fields, the former is adopted.

Visibility. A great advantage of the data used in this paper is the ability to quantify the visibility of research projects (Dai et al., 2022). By having access to the number of downloads for each paper on SSRN, it is possible to approximate the number of people who are likely to have read more than just the abstract of a given paper. There are no natural counterparts to this measure in the patenting literature, as patent offices do not track who accessed which patent over time. Herein lies the major strength of the data at hand, going far beyond what would be possible using patent data. This visibility measure is constructed as the number of downloads a paper has received ($\#downloads$). At the author-year-level, the downloads of all papers first uploaded in year t are averaged ($avg.\ downloads$). As this variable is heavily skewed, the log transformation is used in all analyses ($\ln(1 + avg.\ downloads)$).

Impact. Following the prior literature (Hall et al., 2001), research impact is quantified using the number of forward citations ($\#citations$). Averaging over all papers originally written in year t, an author-year version of the variable is constructed ($avg.\ citations$). Once again, the log is taken to address skewness in this variable ($\ln(1 + avg.\ cites)$). Unlike with patents, a working paper need not end up in a journal to receive citations.

The sample period naturally starts in 2006, with the launch of Twitter. Author-year variables are therefore constructed for starting with the greater of the year of the author's first working paper uploaded to SSRN and 2006. Summary statistics of the final sample are presented in the following section.

3.2.4 Summary Statistics

Before discussing the summary statistics of the final dataset, it is a sensible exercise to assess the selection into using the social network. Intuitively, one would expect that more prominent researchers have greater incentives to set up Twitter accounts, as they anticipate greater followership. To assess the applicability of this notion, Table 3.1 breaks down the sample by authors that are on Twitter and those that are not. Of the economists in the SSRN sample, 13.5% have Twitter accounts. While it is expected that a lot of individuals are falsely identified as not being active on Twitter, the large share of economists who are lends credibility to this setting being used to study interactions between scholars. As individuals in the SSRN sample are expected to differ along many dimensions (see Table 3.2), the set of NBER economists, who are a lot more homogeneous, is considered as well. The share of members who have Twitter accounts is sizable at 44.5% (716 out of 1,608), underpinning the overall relevance of the social network for the profession.

Panel A Table 3.1 is based on the full sample of SSRN authors in the ERN and FEN networks. It must be stated that the identification of Twitter accounts, while sophisticated, is not perfect. Therefore, there naturally are individuals in the non-Twitter subsample who, in fact, have accounts on the platform. However, these individuals are likely significantly less active in research, as well

Table 3.1: Summary Statistics – Authors

This table depicts the summary statistics of the sample of economists. The sample is split based on whether an author could be linked to a Twitter account. While the linking of individuals to Twitter accounts relies on a sophisticated algorithm, there likely are authors in the sample of non-users, which have an account but could not be linked. The sample of Twitter users is thus naturally biased towards researchers that are more active on social media. Panel A presents the summary statistics for all individuals with (at least) one SSRN account. Those authors that could be linked to their Twitter accounts were more prominent along every dimension. To control for differences in tenure and otherwise unobservable ability, the sample is restricted to NBER associates and fellows in Panel B.

	Twitter		**Non-Twitter**		Δ
	Mean	S.D.	Mean	S.D.	$nt - t$
Panel A: SSRN Authors	$N = 19,422$		$N = 124,211$		
SSRN WPs	11.004	17.301	7.521	12.347	-3.483***
NBER WPs	1.211	6.218	0.336	2.933	-0.875***
Papers Total	15.725	26.711	9.792	20.056	-5.934***
Log SSRN Downloads	6.032	2.365	5.628	2.229	-0.403***
Log SSRN Citations	1.812	1.819	1.356	1.557	-0.456***
Panel B: NBER Members	$N = 716$		$N = 892$		
SSRN WPs	33.334	32.541	34.000	29.168	0.666
NBER WPs	22.000	22.634	22.053	21.486	0.053
Papers Total	47.837	52.918	49.719	44.492	1.882
Log SSRN Downloads	7.735	1.863	7.980	1.758	0.246**
Log SSRN Citations	5.037	2.146	5.499	1.886	0.463***

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

as on the social network. Comparing the two groups, scholars on the platform are significantly more productive and generate more impactful research. This is not surprising due to a combination of two effects. Firstly, it is an established fact that the distribution of research output is heavily skewed (Brogaard et al., 2018). Secondly, while having a Twitter account may not be significantly correlated with research output, it is inherently easier to identify the accounts of people stating their job description and affiliation in their biography, which correlates with productivity. As a first pass at assessing whether selection into social media use is driven by ability or tenure, the data is restricted to NBER associates and fellows in Panel B. In the absence of further control variables that could otherwise proxy for ability, the NBER's member selection process ensures much greater homogeneity of the sample.[11] To the extent that research output is a good proxy for ability, this does hold up in the data, as the differences in working papers written are insignificant between the two subsamples.

To get a better sense of the evolution of the network, Figure 3.3 depicts the evolution of the social network in terms of size as well as activity. Figure 3.3a presents user sign-ups as well as the total number of users between 2006 and 2022. While there was a large influx of users around the financial crisis, this pattern is consistent with the overall growth of the platform. From 2014 onwards, quarterly sign-ups have stabilized at around 250 users. To get at the distribution of activity over time, Figure 3.3b depicts the evolution of users for different activity thresholds. The solid red line represents the count of accounts with at least one tweet in a month, whereas the other lines correspond to the number of active accounts when defining activity as more than 5, 10, and 25 monthly tweets, respectively. Several interesting observations can be made. Firstly, there are a lot of passive users, who do not tweet. Secondly, there was a very pronounced activity spike at the onset of the COVID-19 pandemic, especially for highly active accounts.

Moving on to the dataset that is at the core of this paper, the final author-year sample consists of 195,269 author-year observations. Table 3.2 presents the corresponding summary statistics. Panel A is restricted to data from SSRN, while Panel B presents the variables constructed from Twitter. The average academic in the matched sample uploads slightly less than one working paper (0.907) to SSRN each year. However, there is a considerable skew in this productivity measure. Unsurprisingly, a similar pattern emerges in the visibility and impact measures. These observations support the use of logarithmic transformations in the subsequent, quantitative analyses. The average joining year is 2014, corresponding to the half-way point between the platform's launch in 2006 and the end of 2022. This is consistent with the approximately linear growth in the number of economists on Twitter (see Figure 3.3a), which outpaces the relative increase in users of the platform as a whole. A large share of users is silent, as indicated by the small average in $tweeted$ compared to $had_account$. This observation in consistent with Filippas and Horton (2021), who

[11]See `https://www.nber.org/about-nber` for further details on the nomination and selection process.

Figure 3.3: #EconTwitter Users

The figure below depicts the evolution of economists on Twitter since the launch of the platform in 2006. The starting point is the set of economists constructed based on authors of working papers uploaded to the SSRN *Economics Research Network* (ERN) and *Financial Economics Network* (FEN), as well as individuals with profiles on *Research Papers in Economics* (RePEc) or the *National Bureau of Economic Research* (NBER). Individuals are then linked to their respective Twitter accounts.

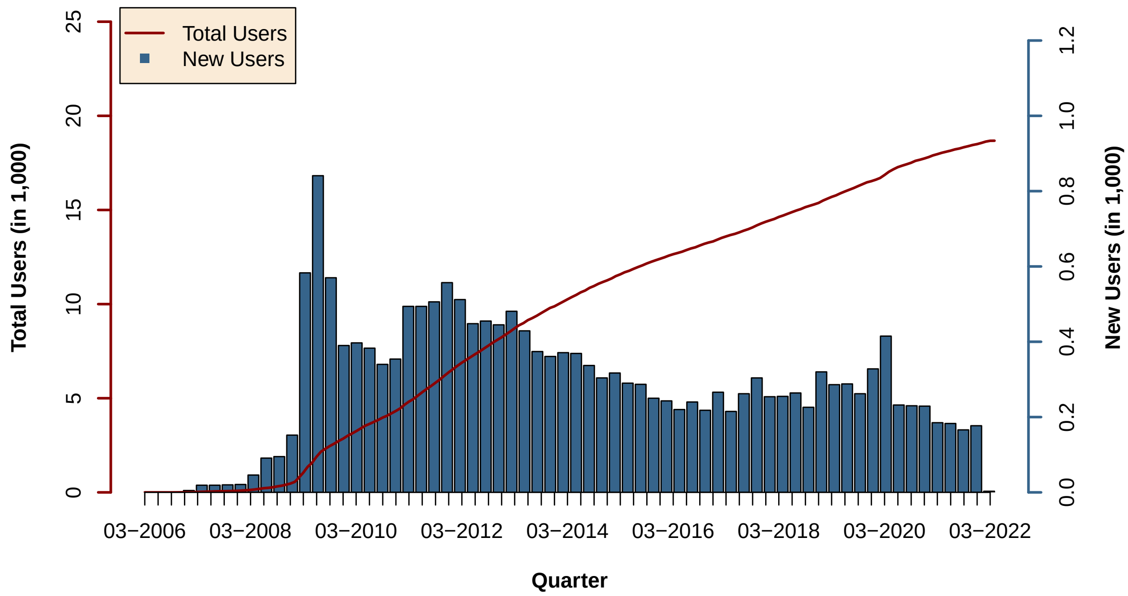

(a) Total Users

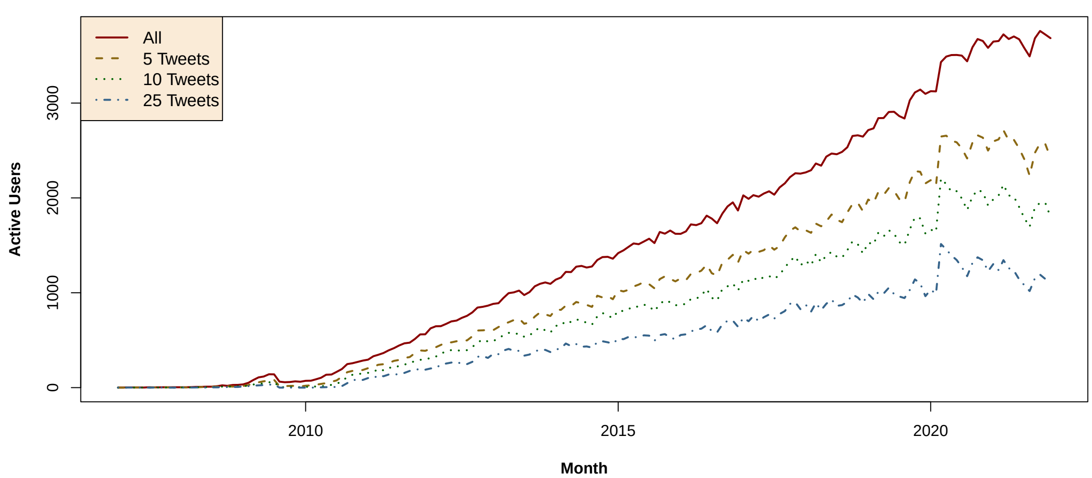

(b) Users by Activity

Table 3.2: Summary Statistics – Cross-Section

This table depicts the summary statistics of the author-year sample used throughout the paper. The sample includes all author-year observations starting after the year the author uploads their first paper to SSRN. All authors in the sample have a Twitter account by the end of the sample period. All variables in panel A are obtained from SSRN's paper and author metadata, aggregated to remove duplicate author accounts. Download and citation variables pertain to the lifetime download/citations gathered by papers written in year t. Variables in panel B are obtained from the Twitter API.

Variable	**N**	**Mean**	**S.D.**	**Min**	**Max**
Panel A: SSRN					
papers	195,269	0.907	2.073	0	159
avg. downloads	195,269	186.201	1,164.511	0	91,482
ln(1+avg. downloads)	195,269	1.622	2.283	0	11.256
avg. citations	195,269	0.125	0.987	0	144
ln(1+avg. cites)	195,269	0.054	0.259	0	4.977
Panel B: Twitter					
joining_year	195,269	2,013.678	3.816	2,006	2,022
had_account	195,269	0.579	0.494	0	1
tweeted	195,269	0.448	0.497	0	1
ln(1+tweets)	195,269	1.920	2.532	0	10.973

similarly find a that silent readers are the dominant members in social networks.

Lastly, Table 3.3 presents the data on daily downloads. These variables are available for SSRN working paper, which have been linked to in a tweet. Panel A presents some of the paper metadata. The average SSRN paper shared on Twitter is recent, with more than half having been written between 2013 and 2019. This already hints at the use of the social network to disseminate recent research, rather than older papers. The number of authors does not differ significantly from the distribution in the sample of other papers written at around the same time. Panel B presents summary statistics of the daily downloads as well as total downloads as of a given day. Most working papers are not downloaded daily, seeing a download on less than one out of every four days.

3.3 Empirical Strategy

The dataset described in the previous section is now used to better understand how interactions shape the research process. As a simplification of reality, research is treated as a linear four-stage

Table 3.3: Summary Statistics – Daily Sample

This table depicts the summary statistics of sample of daily downloads. Panel A contains variables related to the characteristics of the papers in the sample. Panel B presents the summary statistics for the daily downloads.

	N	Mean	SD	p25	Median	p75
Panel A: Paper Details						
# authors	37,305	2.542	2.237	2.000	2.000	3.000
year	37,305	2,015.526	4.476	2,013.000	2,016.000	2,019.000
# downloads (total)	37,305	274.100	2,105.634	21.000	54.000	158.000
Panel B: Daily Data						
# downloads	97,617,803	0.117	3.644	0.000	0.000	0.000
# downloads (cum)	97,617,803	219.674	1,773.615	14.000	40.000	121.000

process.

$$\underbrace{\text{Observations} \rightarrow \text{Problem/Question}}_{\text{Idea Generation}} \rightarrow \text{Research} \rightarrow \text{Outcome} \tag{3.1}$$

Same as every economic agent, researchers observe the world around them. However, driven by their inherent curiosity (Bell et al., 2019), these stimuli lead to the identification of unsolved problems and the formulation of research questions. While in the context of Economics we might want to think of questions such as the effect of quantitative easing on the economy after having observed market movements, this way of identifying problems is shared by innovators more broadly. The creation of many innovations in the private sector similarly goes back to the identification of a problem. Naturally, the first two stages of the process therefore are intertwined and cannot be disentangled. They are thus grouped together and referred to as the *idea generation* stage. Once a viable idea has been identified, the researcher(s) will work on a solution to the problem at hand. This process can be exceptionally long and should be thought of along the lines of Stein (2008). As soon as the innovation has been completed to the researcher's satisfaction, it is turned into a product. Using academia as the setting for this study, work-in-progress research, as reflected in working papers, cleanly falls into the *research* stage, whereas publications are part of the *outcome*. Crucially, as in all models on innovation (Bobtcheff et al., 2016; Hill and Stein, 2021), the innovator can decide to abandon the project at any point in time.

Social interactions at the various stages of the research process may have vastly different implications. Both qualitative as well as quantitative statements will be made, the former to develop a better understanding of the general nature of interactions, while the latter are designed to quantify such effects. On the *qualitative* side, it will be first be assessed what groups of people interact with one another. This is achieved by leveraging the full depth of the data, including information on follower networks, textual data, as well as the link from authors to research projects. On the *quantitative* side, three different, important outcomes measures will be related to the extent of social interactions: (1) research productivity, (2) research visibility, and (3) research impact. The following sections outline how these analyses are structured and to what extent the relationships allow for a causal interpretation. Note that these three aspects map naturally to the research process.

3.3.1 Fixed Effect Model

As a starting point in quantifying the relationship between social interactions and research outcomes, a standard two-way fixed effects model is estimated. The model is specified as follows:

$$y_{i,t} = \beta \times social_{i,t} + \Gamma X_{i,t} + \delta_i + \delta_t + \varepsilon_{i,t} \tag{3.2}$$

The dependent variables $y_{i,t}$ are chosen from those constructed in the previous section. For each author i, they capture research outcomes in year t. β, the coefficient of interest, captures the relationship between social interactions, as measured by $social_{i,t}$, and the dependent variable. $X_{i,t}$ contains (time-varying) controls. δ_i and δ_t are author and time fixed effects respectively, to control for time-invariant ability and a general time-trend in the twitter data. In some specifications, the latter is replaced with an area-year fixed effect, $\delta_{a,t}$, where area is defined as the JEL code of author i. This absorbs differences in general interest levels between fields over time. To give one example, research in financial economics received disproportionately more attention after the great financial crisis. If there are no time-varying unobservables within a given author, the relationship could be interpreted as a causal one. However, it is unlikely that there are no such effects, necessitating a different identification strategy which yields causal estimates.

3.3.2 Event Study

To alleviate concerns regarding reverse causality between social interactions and research outcomes at higher levels of aggregation, time-stamped downloads are used to assess the causal impact of individual tweets on a research project mentioned therein. A window of $[-5, 5]$ is chosen and papers that have been uploaded less than 10 days before the event are removed. The reduced-form regression takes the following form:

$$y_{p,s} = \beta \times post_{p,s} + \delta_p + \delta_t + \varepsilon_{p,t} \tag{3.3}$$

The analysis is conducted at for each paper p, which appears in a tweet and for which daily downloads are observable. Let s denote the days from treatment, i.e., $s \in [-5, 5]$ and $s = 0$ is the treatment period. $post_{p,t}$ is a dummy variable that takes on a value of one if $s \geq 0$. As in (3.2) δ_p and δ_t are paper and time fixed effects. To visualize the effect over time and assess the presence of a pre-event response, a dynamic specification is estimated as well, by replacing $post_{p,s}$ in (3.3) with a series of dummy variables:

$$y_{p,t} = \sum_{s \in \{-5,\ldots,5\}} \beta_s \times D_{t+s} + \delta_p + \varepsilon_{p,t} \tag{3.4}$$

D_{t+s} is an indicator variable for period $t + s$. The coefficients β_s have the natural interpretation of the average of the dependent variable at time s after controlling for the fixed effects.

3.4 Qualitative Results

In this section, the dataset is used to develop a better understanding of the nature, content, and potential drivers of interactions between researchers. Furthermore, the data is used to get a sense of where along the research process these interactions take place (see (3.1) for details).

3.4.1 Information Flow

Considering the growing literature on peer effects (Bailey et al., 2018a,b, 2022) and echo chambers (Cookson et al., 2023; Cookson and Niessner, 2020) in various settings, it is important to assess whether similar effects play a role in research. Of major interest are the characteristics of people who interact and to what extent there is positive assortative matching at play. One key advantage of Twitter compared to other social networks is the directionality of followership. This closely mirrors how we think about interactions in a wide range of settings, where the speaker is not necessarily aware of who their listeners are. While social interactions in the narrow sense are bilateral, in the context of research it is particularly important to make this distinction. Thought leadership plays a key role in many research settings, be they academic or in the private sector, which is easily captured by the selection of individuals into following others. While there might be many things driving the decision to follow a fellow economist on Twitter, the exposure to their ideas likely is the dominant one. Due to the design of the platform, the person being followed oftentimes is not informed of the marginal follower they attracted, limiting the relevance of career concerns in the decision of the follower. Furthermore, since many users, on the platform as well as in the profession, as 'silent readers,' this approach alleviates concerns that active participation in online discussions is endogenous. Unfortunately, the directed network of followership is only available cross-sectionally as of February 2022, making it impossible to analyze any time-series variation.

As with any analysis pertaining to the interaction between individuals, the overarching question of whether there is homophily at play comes to mind. In the case of the research process, there are many reasons for why interactions with scientists from both narrowly (homophily) as well as more distantly related fields are valuable. On the one hand, interactions with individuals working in the same area can be valuable as these close peers are ultimately the target audience of the research outcome. This naturally makes their comments more relevant, which is why rational agents should put greater weight on their views. A countervailing force is the increase in the risk of being scooped (Hill and Stein, 2020), giving rise to competing papers, reducing these incentives. On the other hand, it oftentimes requires outside-the-box thinking to solve problems, which scholars from less related fields might bring. One, anecdotal, example is that of Black and Scholes (1973), who were only able to solve their model by leveraging insights from a befriended Physicist. It therefore is not ex-ante clear which of the two effects should dominate.

Using data on the authors' research areas (as defined by JEL codes) and followers, the directed network of followership is used to construct concentration measures of followership between economic subfields. Let the network be such that each researcher is an edge, which is linked to other edges through vertices. Each vertex originates from the follower and ends at the one being followed. Let $n\_follower_{d \to e}$ be the number of vertices from research area d to e in the network. The share of followers of research area e that originates from area d is then computed as

$$follower\_share_{d \to e} = \frac{\#followers_{d \to e}}{\sum_{d'} \#followers_{d' \to e}}.$$

This measure must be read as the share of all followers of e belonging to group d and can be thought of as the importance of the in-group relative to the out-group in research. Naturally $\sum_d follower\_share_{d \to e} = 1$, as this measure is constructed over the incoming vertices.

Figure 3.4 depicts the variable $follower\_share_{d \to e}$ for all values of d (row) and e (column). In the interest of brevity, only the most striking observations will be highlighted. Consistent with homophily, the diagonal is very pronounced, with most followers of a given area themselves belonging to that area. As seminars and conferences are usually area-specific, it must be noted that such connections are naturally easier to form as well. Comparing different subfields, JEL code G (Financial Economics) stands out, as there is a much stronger clustering of followers within the same group. While being an important field within economics, Financial Economics naturally stands secluded due to a large fraction of scholars being located at business schools and thereby spatially separated from their colleagues in economics departments. Furthermore, while financial economists routinely publish in some of the top journals in economics, the reverse is not as common. In stark contrast, JEL codes D (Microeconomics) and J (Labor Economics) have a much

Figure 3.4: Who Listens to Whom?

The figure below depicts the proportion of followings between economists in different research areas based on JEL codes. Each cell (r, c) represents the share of area c's followership originating from area r. To construct the measure, economists with Twitter accounts are assigned JEL codes based on their publication records. The set of followers for all accounts is current as of June 2022 ($N = 9,611$). Only followers that were assigned a JEL code are considered. See Table C.1 in the Appendix for a mapping between JEL codes and research areas.

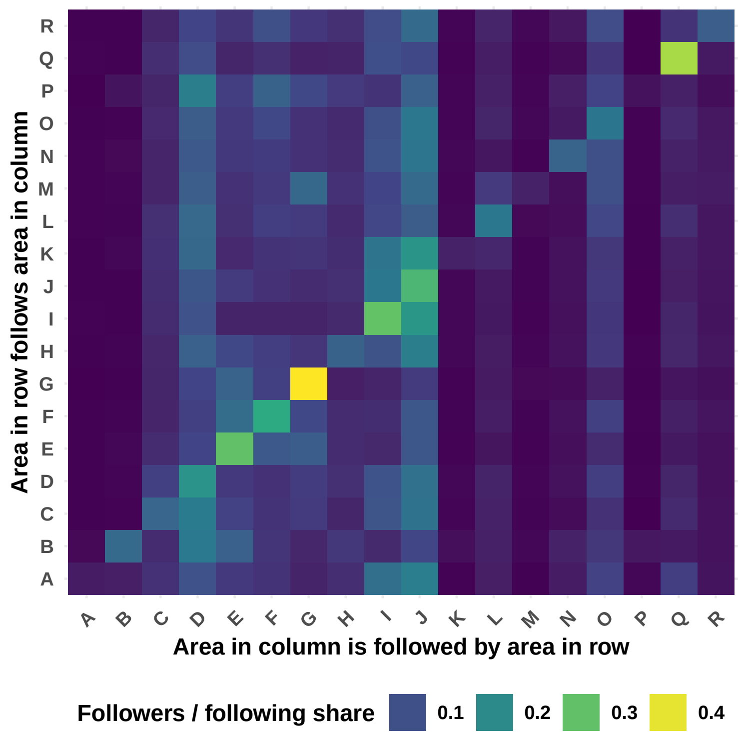

more evenly distributed set of followers. In the case of Microeconomics this likely is driven by the relevance of the field as a building block of both theoretical as well as empirical work in other fields. Labor Economics on the other hand has always been a very interdisciplinary area, which can explain this phenomenon.

It can be concluded that interactions do primarily occur within the same research area. This indicates that information primarily flows primarily to in-group members, although there is a lot of heterogeneity across fields. These significant differences with respect to 'openness' to other fields are striking and may help in explaining why some fields have grown more rapidly than others over

Figure 3.5: Links to Outside Websites

All tweets by economists in the sample are searched for links to outside websites. These are then aggregated to the corresponding domain. News websites are depicted in blue, websites related to academic research in red, and others in gray.

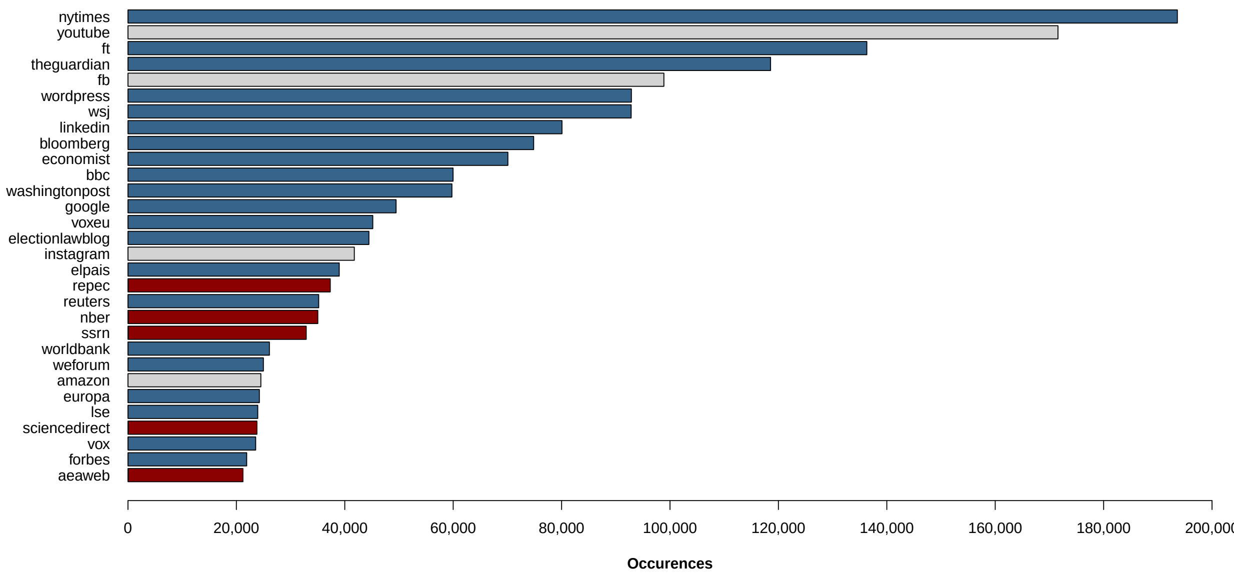

the past decade.

3.4.2 Information Sources

Moving from the network to the information shared, links to outside web pages are extracted from the sample of tweets. These are useful in assessing the type of interaction taking place on Twitter and the kinds of information disseminated. As an example, if a tweet includes a link to a working paper on SSRN, the tweet is classified as having referred to SSRN and linked to the exact paper mentioned in it. This process necessitates the assignment of individual URLs to their corresponding domains (e.g., nytimes.com). The relative frequency of the most shared websites is depicted in Figure 3.5. URLs that are associated with news stories are colored in blue, whereas those that pertain to research papers have been highlighted in red. There is a large number of domains that fall into neither of the two, which are colored grey.

The most striking observation is that websites related to (economic) news are shared (blue bars) most frequently. One possible interpretation is that, as one would expect from economists, individuals follow the news leading to the validation or falsification of existing research as well as the inspiration for new projects. Arguing along these lines, these types of interactions broadly fall into the idea generation stage. That being said, there obviously are many other potential mechanisms at play, which cannot be ruled out.

Directing attention to those websites that relate to research more concretely (red bars), there seem to be two distinct groups of websites being shared. The domains most linked to are RePEc, the NBER, and SSRN. These sites have in common that they serve as sources of pre-prints rather than publications. Outlets of published work (e.g., ScienceDirect, AEAweb) appear much less frequently. While this is consistent with interaction being a lot more relevant when working papers are still in the process of being revised, it must be acknowledged that there naturally are far more pre-prints than publications. There are two potential stories at play here, which are not mutually exclusive. On the one hand, for a paper to have a high impact, it must be marketed within the profession. Twitter naturally is one potential way of achieving this goal (as a complement to presenting at conferences). On the other hand, it is also at this stage that comments from other economists are most valuable, as they can be incorporated into the next revision, whereas published papers are set in stone. While both arguments are compelling, the obvious concern remains that there simply is a much larger number of working papers compared to published work, which might mechanically drive these results.

Lastly, a substantial number of websites that are being shared (those in gray) correspond to neither news, nor research. Some of these references might still relate to economic research, though it is much harder to deduce this relationship from the domain alone. This observation definitely warrants a more thorough investigation of the informational content of the corpus in this paper, leveraging novel NLP techniques. Whatever the case may be, the prevalence of links to Google, YouTube, and Co. underpins the relevance of defining #EconTwitter as interactions that are related to economics, as there otherwise would be many tweets considered to be related to research, when they clearly are not.

Taken at face value, the implications of this finding are that social interactions decline along the research process. To the extent that there are many more potential ideas to be worked on than viable ones (Bennett and Robinson, 2022), this is consistent with a funnel. Whereas it takes many interactions with others to zero-in on a few feasible projects, once these have been identified, the risks of discussing (i.e., being scooped) start to outweigh the benefits. Once the outcome stage has been reached, further input is no longer valuable. This therefore validates the choice of research measures constructed in the previous section, focusing on both the marketing of research projects (visibility), as well as their ultimate impact. To shed some more light on the potential motivation behind sharing research papers, the following section will use information of the exact papers shared to assess whether promoting one's own research or endorsing that of other is the key driver of these findings.

3.4.3 Endorsement and Promotion

As outlined in the previous section, a not insignificant share of conversations revolves around in-progress as well as completed research projects. Leveraging the granularity of the data, each tweet that contains a reference to an academic paper is linked to said working paper or publication. This affords many things, primarily, it allows for the identification of whether the person sharing the paper is one of the coauthors. The variable $self\_share$ is a dummy variable that takes on a value of one if the paper is shares by one of the authors and zero otherwise. At the author-level, the share of self-references is then computed as:

$$self\_share_i = \frac{\sum_p self\_ref_{ip}}{\sum_p \mathbb{1}\{self\_ref_{ip} \geq 0\}}.$$

These variables are constructed over the set of tweets which contain a reference to an academic paper. They can thus be interpreted as the empirical equivalent of the probability of sharing one's own research conditional on writing a tweet that references a paper. The two extreme values of the measure, i.e., 0 and 1, will be referred to as *endorsement* and *promotion*. The motives behind

Figure 3.6: Endorsement and Promotion

For each Twitter account the number of links to (working) papers (co-)authored by the account holder is divided by the total number of papers shared. The resulting ratio ($self\_share_i$) is bound between zero and one, and captures the extent to which the individual uses Twitter to promote their own work ($self\_share_i = 1$), or endorse others ($self\_share_i = 0$).

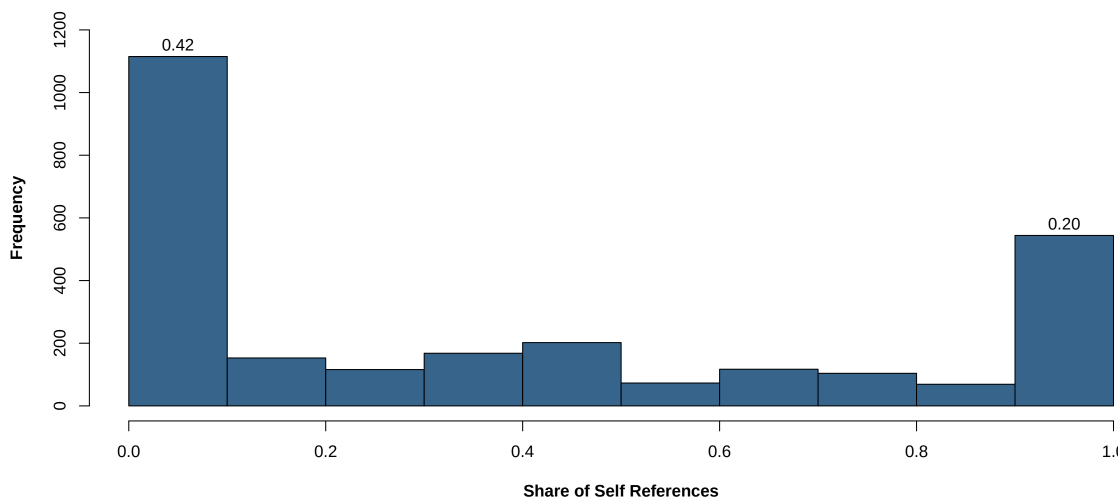

sharing research from others likely are quite different from those compelling an author to share their own work.

Equipped with this measure, it is natural to ask to what extent either extreme prevails in the data. This is equivalent to asking if researchers are actively seeking exposure of their work to comments from their peers or whether motives of endorsement are the prevailing force (Garicano and Santos, 2004). One can differentiate between the promotion of a paper by one of the authors and the endorsement by an unaffiliated individual. Figure 3.6 plots the empirical distribution of $self\_share$. It is clear from the plot that there is a lot of heterogeneity in terms of whether authors use Twitter to only promote their own work or endorse papers by other individuals. While there is greater mass on $self\_share < 0.5$ than $self\_share > 0.5$, there seems to be a fair bit of heterogeneity. There furthermore are stark differences between subfields. Financial economics seems to exhibit the reverse, i.e., the promotion of an author's papers dominates is more prevalent than the endorsement of the work of others.

While there are many reasons for why someone may want to engage in either behavior, it is not clear what the benefits of doing so are. Either behavior can be in the interest of the individual, as being perceived as someone knows the literature well can have positive career effects. Similarly, while sharing one's own research can have positive effects in the short-term, there may be backlash when doing so exclusively. To what extent these two correlate with research outcomes will be analyzed in the following section. Furthermore, it is not clear to what extent writing about a paper on Twitter has any measurable impact in the first place. In an event study, this effect will be quantified towards the end of the following section.

3.5 Quantitative Results

In more quantitative analyses, proxies for interactions among individuals are related to the research outcomes as defined in Section 3.2. There are three variables of interest are (1) research productivity, (2) research visibility, and (3) research impact. To quantify the relationship, the specifications outlined in Section 3.3 are estimated.

3.5.1 Productivity

Following the research process as outlined in (3.1), the first outcome of interest is productivity, i.e., the number of projects an individual starts in a given year. Considering a model of the form as in (3.2), the sign of the coefficient of interest β is ex-ante ambiguous. On the one hand, interactions are naturally correlated with the size of one's network, which one would expect to have a positive effect on productivity. On the other hand, to the extent that researchers are time constrained, any increase in social interactions will lead to a reduction in time spent on research, thereby reducing productivity on the margin.

Coefficient estimates from regressing the number of working papers uploaded to SSRN, as a measure of research productivity, on different measures of social interactions are reported in Table 3.4. In columns 1 and 3, the sample includes all authors in the dataset, whereas inactive

Table 3.4: Fixed Effect Model – Productivity

The table below reports point estimates from regressing the number of papers uploaded to SSRN by an individual in a given year on a dummy variable whether that individual tweeted in the same period ($tweeted$), as well as the log of the number of tweets in the same period ($\ln(1 + tweets)$). In columns 1 and 3, the sample includes all authors, irrespective of the presence of a twitter account. Columns 2 and 4 only include data from accounts that have been linked to Twitter accounts. All regressions include fixed effects for the author and the area-year, where area is defined based on the author's JEL code. Standard errors are robust and clustered at the author-level.

Dependent Variable:	# ssrn papers			
Model:	(1)	(2)	(3)	(4)
$tweeted$	0.0499***	0.1090***		
	(0.0159)	(0.0193)		
$\ln(1 + \#tweets)$			0.0091**	0.0282***
			(0.0036)	(0.0039)
Author FE	Yes	Yes	Yes	Yes
Area-Year FE	Yes	Yes	Yes	Yes
Sample	all	active	all	active
Observations	195,269	113,108	195,269	113,108
R^2	0.346	0.443	0.346	0.443

Clustered (author) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

accounts are removed from the sample in columns 2 and 4. All regressions include fixed effects for the author and area-year, where area is defined based on the author's JEL code. Standard errors are robust and clustered at the author-level. In columns 2 and 3, the independent variable $tweeted$ is an indicator of whether the author wrote at least one tweet. This measure of the extensive margin should not be interpreted as having engaged with peers in general. Rather, it lends itself to a more natural interpretation of having been active on Twitter in the loosest sense. Individuals for whom the dummy switches on are those that presumably have had at least some exposure to the discussions on #EconTwitter. To understand the actual relationship between network effects and productivity more cleanly, columns 3 and 4 focus on the intensive margin. Replacing the independent variables with the logarithm of the number of tweets written in a given year ($\ln(1 + \#tweets)$), the coefficient measures that marginal effect of an increase in the number of tweets on research productivity.

Focusing on the extensive margin results in the first two columns, the relationship is positive and highly significant. In moving from column 1 to column 2, inactive accounts are removed from the sample. Based on the preferred specification in the second column, individuals who wrote at least one tweet produced about 0.1 additional working papers in the same year. This is an increase of about 11% of the unconditional mean, which is not just statistically significant but also economically meaningful. While the inclusion of author and area-year fixed effects absorbs some of the heterogeneity in skill, it cannot alleviate the concern that reverse causality is driving these results. In fact, this likely is the case, as individuals who authored a paper are disproportionately more likely to tweet about it in a 'new working paper alert' following the upload to SSRN. In an untabulated robustness check, in which tweets relating to own research are removed from the measure of social interactions, the coefficient remains significant and only suffers a small attenuation, alleviating some of these concerns.

Moving on to the intensive margin results in columns 2 and 4, the relationship once again is positive and highly significant. In the preferred specification in column 4, a one standard deviation increase in the independent variable is associated with an increase of about 0.07 in the number of papers written. This effect, which is about 30% smaller than the extensive margin result, is equally statistically significant and remains economically meaningful. To the extent that the measure of interactions on the social network is a proxy for the time an individual spends 'socializing' or floating ideas with peers, these results are consistent with interactions being helpful in accelerating the research process. By considering the number of working papers as the dependent variable, the results specifically speak to the value of interactions when moving from the *idea generation* to the *research* stage. The positive effect is consistent with the intuition outlined previously. As before, removing tweets related to an individual's own research does not lead to material changes in the conclusion.

In summary, across all specifications, a positive association between interactions on Twitter and research output can be identified. Unsurprisingly, when moving from the sample of all economists to those that have active accounts, the effects become stronger and more significant. Nevertheless, concerns regarding reverse causality remain significant. While removing tweets promoting an individual's own research from the sample do not materially change the results, this does not address the concern in its entirety. The inclusion of author fixed effects ensures that the identification of the coefficients comes from variation within an author rather than across authors. Nevertheless, the work of more productive individuals is more likely to appear in online discourse, even if they are not the ones writing about it. This can then incentivize them to contribute, i.e., lead to reverse causality. Due to the nature of research and social networks, absent experimental settings, it unfortunately is difficult to alleviate such concerns.

3.5.2 Visibility

The second research outcome of interest which will be examined next is the visibility of working papers. While visibility is impossible to quantify in private sector innovation, by focusing on academic research in Economics and using data from SSRN, it becomes tractable. Firstly, since pre-prints play a key role in academia, research results become observable in real time, whereas even patent filings in the U.S. are only disclosed after 18 months. Secondly, using data on the number of downloads of a working paper from SSRN, the potential number of readers can be quantified much more cleanly (Dai et al., 2022). This is in stark contrast to the existing literature in innovation, which relies heavily on patents. Even after a patent is published, it is impossible to assess the number of individuals who were exposed to it. This measure therefore does not have an analogue in the patent literature. While SSRN is not the only way of accessing working papers, it is the dominant one and can serve as a meaningful proxy for the number of 'eyeballs' a pre-print has gathered.

The model, which will be estimated to quantify the relationship between social media activity and downloads is described in (3.2). Naturally, the promotion of one's research should have a direct effect on the number of downloads a paper gathers over the course of its existence. An insignificant or even negative relationship would defeat the purpose many scholars see in engaging with one another. Anything other than $\beta > 0$ would therefore be a very surprising result. As in

Table 3.5: Fixed Effect Model – Visibility

The table below reports point estimates from regressing the log of the number of downloads of papers written by author i in year t from SSRN ($ln(1 + avg.\ downlaods)$) on a dummy variable whether that individual tweeted in the same period ($tweeted$), as well as the log of the number of tweets in the same period ($\ln(1 + tweets)$). In columns 1 and 3, the sample includes all authors, irrespective of the presence of a twitter account. Columns 2 and 4 only include data from accounts that have been linked to Twitter accounts. All regressions include fixed effects for the author and the area-year, where area is defined based on the author's JEL code. Standard errors are robust and clustered at the author-level.

Dependent Variable:	ln(1+avg. downloads)			
Model:	(1)	(2)	(3)	(4)
$tweeted$	0.0385***	0.0707***		
	(0.0149)	(0.0210)		
$\ln(1 + \#tweets)$			0.0108***	0.0208***
			(0.0033)	(0.0047)
Author FE	Yes	Yes	Yes	Yes
Area-Year FE	Yes	Yes	Yes	Yes
Sample	all	active	all	active
Observations	195,269	113,108	195,269	113,108
R^2	0.371	0.434	0.371	0.434

Clustered (author) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

the previous section, two measures of social interactions are going to be considered, whether the individual was active on the platform ($tweeted$), and (2) the log-transformation of the number of tweets in a given year ($\ln(1 + \#tweets)$).

Coefficient estimated are reported in Table 3.5. As in the previous section, the full sample is considered in columns 1 and 3, whereas inactive accounts have been removed from the sample in columns 2 and 4. The same author and area-year fixed effects are included in all specifications, while standard errors remain clustered at the author-level. Starting from the high-level results, the coefficient estimates are all positive and highly significant, as one would expect them to be. As before, restricting the sample to active accounts leads to a strengthening of the effect.

The extensive margin results in columns 1 and 3 suggest a significant impact of social media presence on the visibility of the author's research. In author-years with activity on twitter, the average number of downloads increases by between 3.85% and 7.07%. The relationship is highly significant and economically meaningful. While it remains to be seen to what extent the increased visibility translates into a greater number of citations, the result suggests the presence of tangible benefits of this kind of social interaction. While the identification is within-author, due to the presence of an author fixed effect, there are two key channels it cannot address. Firstly, authors may select into using Twitter in exactly those years when they have viable papers that they want to advertise. While the removal of tweets that mention the author's own papers addresses this somewhat, it cannot entirely alleviate the concern. Secondly, the outcome measure is the average number of downloads of an author's papers which were written in year t. One could argue that it would be more meaningful to consider the number of downloads all their papers have received in a given year instead. Unfortunately, the data available for the full sample only contains the cumulative downloads at the time of retrieval, omitting any such kind of time-series information.

In columns 2 and 4, the independent variable is replaced with the logarithm of the number of tweets. Since this transforms the specification into a log-log-regression, the coefficient β is to be interpreted as an elasticity. While the coefficients remain positive and statistically significant, their economic significance is slightly diminished. In the preferred specification in column 4, the results suggest that a one standard deviation increase in the number of tweets leads to a 5% increase in the number of downloads. The same concerns regarding the causality of these results still applies. Nevertheless, they are consistent with a significant benefit of an author's social connectedness on the visibility of their research projects.

It can be summarized that there are tangible benefits to an author being actively involved in the profession. Social connections are strongly associated with increased visibility of one's research, though reverse causality cannot be ruled out. Furthermore, there are several potential channels at play. On the one hand, the results may be entirely driven by authors writing about their own research in an attempt at achieving greater numbers of downloads. On the other hand, the increased visibility of the author may just have a halo effect on their research. While the latter is difficult to test empirically, the former, which is the more relevant one, can be assessed from a causal standpoint quite easily. To zoom in on the direct effect of sharing research papers on visibility, an

Table 3.6: Event Study – Visibility

The table below reports point estimates from an event study on the (log) count of daily downloads around a paper being mentioned on Twitter. All regressions include fixed effects for the author and the paper. Standard errors are robust and clustered at the author-level.

Dependent Variables:	$cum.\ downloads$		$\ln(1 + cum.\ downloads)$	
Model:	(1)	(2)	(3)	(4)
$post$	38.74***		0.2283***	
	(3.361)		(0.0233)	
$self\_tweet \times post = 0$		282.0		0.0487
		(172.0)		(0.0417)
$self\_tweet \times post = 1$		299.3*		0.1565***
		(171.7)		(0.0413)
Author FE	Yes	Yes	Yes	Yes
Paper FE	Yes	Yes	Yes	Yes
Observations	142,629	142,629	142,629	142,629
R^2	0.84924	0.84926	0.93344	0.93000

Clustered (author) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

event study using time-stamped download data is conducted. This is done in a reduced-form (3.3), as well as in a dynamic (3.4) way.

Having access to novel data on the daily number of downloads for a subset of articles uploaded to SSRN, a standard event study framework allows for the high-frequency identification of the causal effect of research being mentioned on Twitter on the actual number of downloads. Estimates for the reduced-form event study specification defined in (3.3) are reported in Table 3.6. In columns 1 and 2 the dependent variable is the raw number of cumulative downloads, whereas it is replaced by the log-transformation thereof in columns 3 and 4. Columns 1 and 3 focus on the impact of a paper being mentioned on the downloads, regardless of who brought up a given paper. To assess whether the effect is stronger for promotions or endorsements, the $post$ variable is interacted with a dummy $self\_tweet$, which takes on a value of one if the focal tweet was written by one of the authors and zero otherwise.

The point estimates suggest an increase in the number of downloads by 38.74 or about 14% of the unconditional mean of total downloads. This effect is large and very meaningful, as there likely

Figure 3.7: Event Study – Visibility

The figure below depicts point estimates from an event study of the log count of daily downloads around a paper being mentioned on Twitter. Regressions are estimated separately for tweets that are written by one of the authors of a given paper (Promotion) and those by other individuals (Endorsement). All regressions include fixed effects for the author and the paper. Standard errors are robust and clustered at the author-level.

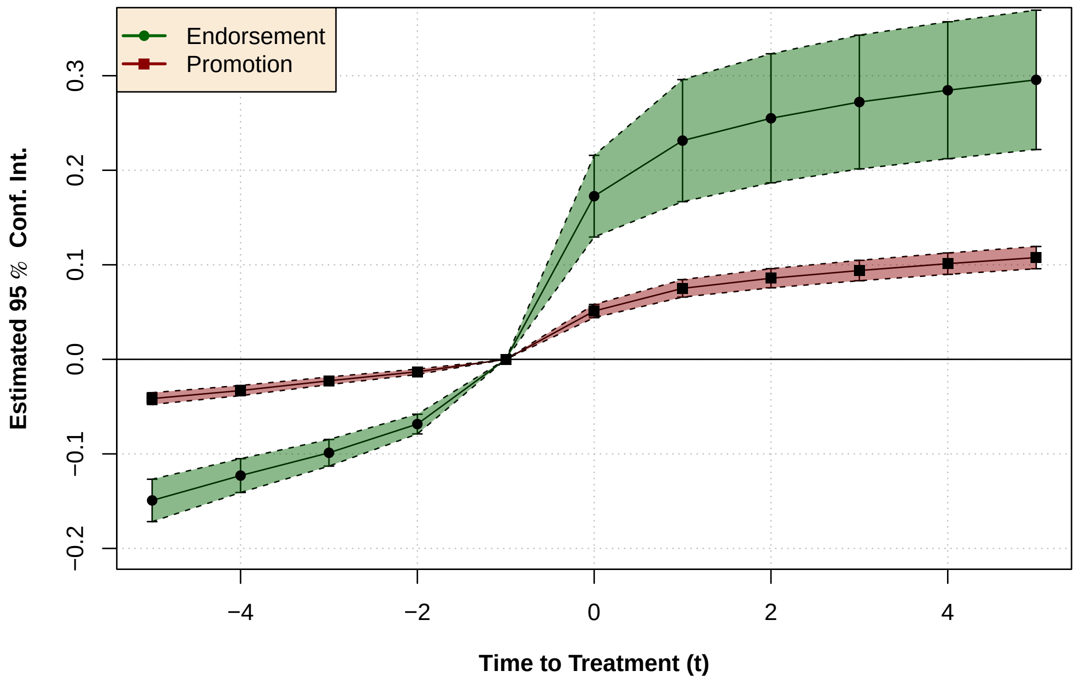

are long-lasting network effects from an early increase in visibility. Focusing on the $t \in [-5, 5]$ window, the effect is similarly meaningful. In column 3 the estimate suggests an increase in downloads by 22.98% over the pre-period. Since these results are derived from a simple event study, concerns remain regarding the dynamics of this effect. Therefore, a dynamic specification (3.4) is estimated and the corresponding coefficients plotted in Figure 3.7. It is of upmost importance not to confuse the figure for a difference-in-differences plot. The two time-series are separately estimated from separate samples and just plotted in the same graph to allow for a comparison. Furthermore, since the dependent variable is the logarithm of cumulative downloads, it is natural that the time-series are trending.

Papers that are being endorsed seem to be different from those that are being promoted, as their pre-event slope is steeper. Nevertheless, it does seem like the impact of an endorsement is greater.

At least within the $t \in [-5, 5]$ window around the event, there does not seem to be evidence of reversal, as the evolution of downloads grows at roughly the same long-run rate after the immediate effect has worn off. The difference in effect size is consistent with endorsements being of greater value, while peers discount the promotion of a given paper.

Summarizing, it can be said that there is a significant impact of social media presence on research visibility. This, first pass, result is unsurprising, as it seemingly is this expectation driving individuals into using Twitter in the first place. The more meaningful result is the exact quantification of this effect. Additional visibility added through sharing has a significant impact. While positive across specifications, there is evidence of differences in the effect size depending on who promotes/endorses a paper, warranting further research in this area. To the extent that there is a trade-off between promoting one's work and sharpening results, it is conceivable that the increase in visibility comes at the cost of research impact. It is this relationship that will be analyzed in the last section.

3.5.3 Impact

So far it seems like interactions do not have any cost. However, as discussed in the previous section, if innovators are time constrained, investing valuable resources in engaging with peers, while helpful in the early stages of research projects (see Table 3.4), could come at the cost of overall quality of research projects. In this section, this potential trade-off will be explored further by considering the average number of citations gathered by papers written in year t as the dependent variable. If there are such trade-offs, one would expect the coefficient on the measures of social interaction to be negative. In the case of complementary, the coefficient should instead be positive in sign. In case of no interaction between the two margins, e.g., when the time constraint is not binding, estimates should be approximately equal to zero.

Coefficient estimates for (3.2), when replacing the dependent variable with the average number of citations gathered by papers written by author i in year t, are reported in Table 3.7. Across all specifications the effects are negative, consistent with the notion of a trade-off between social

Table 3.7: Fixed Effect Model – Impact

The table below reports point estimates from regressing the log of the number of citations to papers written by author i in year t ($ln(1+avg.citations)$) on a dummy variable whether that individual tweeted in the same period ($tweeted$), as well as the log of the number of tweets in the same period ($\ln(1+tweets)$). In columns 1 and 3, the sample includes all authors, irrespective of the presence of a twitter account. Columns 2 and 4 only include data from accounts that have been linked to Twitter accounts. All regressions include fixed effects for the author and the area-year, where area is defined based on the author's JEL code. Standard errors are robust and clustered at the author-level.

Dependent Variable:	ln(1+avg. citations)			
Model:	(1)	(2)	(3)	(4)
$tweeted$	-0.0040**	-0.0013		
	(0.0019)	(0.0024)		
$\ln(1+\#tweets)$			-0.0015***	-0.0011*
			(0.0004)	(0.0005)
Author FE	Yes	Yes	Yes	Yes
Area-Year FE	Yes	Yes	Yes	Yes
Sample	all	active	all	active
Observations	195,269	113,108	195,269	113,108
R^2	0.22000	0.25020	0.22005	0.25024

Clustered (author) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

interactions and research impact. In columns 1 and 2, the presence of an individual on Twitter is used as measure of social connectedness, while columns 3 and 4 focus on the intensive margin, i.e., the number of tweets written. Moving from column 1 to column 2 (column 3 to column 4), inactive authors are removed from the sample.

The extensive margin results in columns 1 and 2 suggest a decrease in citations by up to 0.4%. While negative and statistically significant, this effect is economically negligible. In fact, when restricting the sample to active users, it vanishes, substantiating this notion. Moving to the intensive margin results in columns 3 and 4, the effect of writing a greater number of tweets in a given year is related to the forward citations received by papers written over the same period. These effects are also small, translating into a reduction in citations by less than 1% for a one standard deviation increase in social activity. As in previous regressions, these results do not necessarily allow for a causal interpretation, as there may be time-varying author effects impacting activity as well as output. Nevertheless, the results do shed light on the potential dark side of social interactions in the research process. While certainly not the main reason for the secrecy surrounding innovation in the private sector, these concerns might add to it.

In summary, while their economic magnitudes are small, this result does shed some light on how social interactions impact the research process for the better (visibility) or the worse (long-term impact). This is the expected result if researchers are time-constrained and need to allocate effort between work, which increases output quality, and social interactions that can lead to greater visibility. However, if researchers were rational and aware of the constraints they face, they should allocate their time optimally. From an observer's point of view, it seems that this is not the case, raising further questions as to what their objective function looks like.

3.6 Conclusion

The interactions between innovators and how they shape the innovation process have been of great interest to research in Economics. In this paper, data on the activity of research economists on Twitter (#EconTwitter) is used to take a first cut at the question how specific types of interactions between academics impact the research process. Linking authors to both their publication records, including working papers, as well as their Twitter activity, this is the first paper to use big data from a social network to link interpersonal interactions to innovation outcomes.

As commonly observable in social networks, relationships between researchers show a large degree of homophily in terms of research areas. The degree to which fields favor the in-group differs and is greatest for financial economics. Social interactions center around, in order of importance, news, ongoing research, and completed research. This observation speaks to the relevance of interactions at different points in the research process, consistent with their importance declining as projects mature. Lastly, there is early evidence of two distinct use cases for interactions as they relate to ongoing/completed research projects. Specifically, individuals self-promote their own work as well as endorse that of others. While there are some that engage in both, most individuals cluster at the extremes. The relative importance of these motives differs between fields.

Informed by these observations, proxies for interactions among individuals are related to research outcomes (productivity, visibility, and impact) in more quantitative, reduced-form analyses. Evidence suggests that an increase in social interactions with fellow academics increases research output, as measured by working papers written. However, as the effect is contemporaneous, it is difficult to rule out reverse-causality. With respect to research impact, two distinct measures are considered. While activity in the academic community is positively related to the visibility of research, as measures by downloads of a paper, it negatively relates to the more common citation measure. While not to be interpreted causally, these observations are consistent with a trade-off between improving research projects and marketing them when the authors' time is constrained.

In separate analyses, the causal effect of working papers appearing on the social network on the number of downloads they receive is assessed. Using an event study an increase in lifetime downloads of more than 20% is documented. This finding suggests the importance of visibility of research, especially early on in the roadshow. Splitting the sample into cases of (self-) promotion and endorsement, it is found that this effect is stronger when individuals other than the authors reference a working paper. This is consistent with readers being sophisticated and controlling for the authors incentives when sharing their work. When shared by a third party, these concerns are less relevant, leading to a stronger effect.

This paper is the first to relate interactions between academics to research outcomes at a micro-level. It provides a setting in which not only the occurrence of an interaction can be observed, but also the content of messages exchanged. While not capturing every message sent between individuals, it represents a significant improvement in measurement of interactions. Preliminary results suggest that these have profound effects on research outcomes. Future research should be focused at tightening these results and establishing causality.

Bibliography

Aghion, P., Akcigit, U., Hyytinen, A., and Toivanen, O. (2017). The social origins of inventors. Working Paper, National Bureau of Economic Research.

Aghion, P., Dewatripont, M., and Stein, J. C. (2008). Academic freedom, private-sector focus, and the process of innovation. *The RAND Journal of Economics*, 39(3):617–635.

Agrawal, A. and Goldfarb, A. (2008). Restructuring research: Communication costs and the democratization of university innovation. *American Economic Review*, 98(4):1578–1590.

Allayannis, G., Brown, G. W., and Klapper, L. F. (2003). Capital structure and financial risk: Evidence from foreign debt use in east asia. *The Journal of Finance*, 58(6):2667–2710.

Almeida, H., Fos, V., and Kronlund, M. (2016). The real effects of share repurchases. *Journal of Financial Economics*, 119(1):168–185.

Angrist, J. D. (2014). The perils of peer effects. *Labour Economics*, 30:98–108.

Anton, M., Ederer, F., Gine, M., and Schmalz, M. C. (2021). Innovation: The bright side of common ownership? *Working Paper*.

Antràs, P. and Caballero, R. J. (2009). Trade and capital flows: A financial frictions perspective. *Journal of Political Economy*, 117(4):701–744.

Antràs, P. and Foley, C. F. (2015). Poultry in motion: A study of international trade finance practices. *Journal of Political Economy*, 123(4):853–901.

Arora, A., Belenzon, S., and Sheer, L. (2021). Knowledge spillovers and corporate investment in scientific research. *American Economic Review*, 111(3):871–898.

Atkin, D., Chen, M. K., and Popov, A. (2022). The returns to face-to-face interactions: Knowledge spillovers in silicon valley. Working Paper, National Bureau of Economic Research.

Azoulay, P., Zivin, J. S. G., and Manso, G. (2011). Incentives and creativity: evidence from the academic life sciences. *The RAND Journal of Economics*, 42(3):527–554.

Azoulay, P., Zivin, J. S. G., and Wang, J. (2010). Superstar extinction. *Quarterly Journal of Economics*, 125(2):549–589.

Babina, T., He, A. X., Howell, S., Perlman, E. R., and Staudt, J. (2020). The color of money: Federal vs. industry funding of university research. Working Paper, National Bureau of Economic Research.

Backus, M. (2020). Why is productivity correlated with competition? *Econometrica*, 88(6):2415–2444.

Baersch, S.-E., Heckemeyer, J. H., and Olbert, M. (2018). Transfer pricing and the decision-making authority of the tax function in multinational companies. *Working Paper*.

Bailey, M., Cao, R., Kuchler, T., and Stroebel, J. (2018a). The economic effects of social networks: Evidence from the housing market. *Journal of Political Economy*, 126(6):2224–2276.

Bailey, M., Dávila, E., Kuchler, T., and Stroebel, J. (2018b). House price beliefs and mortgage leverage choice. *The Review of Economic Studies*, 86(6):2403–2452.

Bailey, M., Johnston, D., Kuchler, T., Stroebel, J., and Wong, A. (2022). Peer effects in product adoption. *American Economic Journal: Applied Economics*, 14(3):488–526.

Bank for International Settlements (2024). Debt securities statistics, bis ws_na_sec_dss 1.0. Accessed on 21 February 2024.

Bartov, E., Givoly, D., and Hayn, C. (2002). The rewards to meeting or beating earnings expectations. *Journal of Accounting and Economics*, 33(2):173–204.

Bell, A., Chetty, R., Jaravel, X., Petkova, N., and Reenen, J. V. (2019). Who becomes an inventor in america? the importance of exposure to innovation. *The Quarterly Journal of Economics*, 134(2):647–713.

Benmelech, E., Kumar, N., and Rajan, R. (2022). The secured credit premium and the issuance of secured debt. *Journal of Financial Economics*, 146(1):143–171.

Bennett, V. M. and Robinson, D. T. (2022). Why aren't there more minority entrepreneurs? *Working Paper*.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile Prices in Market Equilibrium. *Econometrica*, 63(4):841–890.

Bertrand, M., Duflo, E., and Mullainathan, S. (2004). How Much Should We Trust Differences-In-Differences Estimates?*. *The Quarterly Journal of Economics*, 119(1):249–275.

Bhojraj, S., Hribar, P., Picconi, M., and McINNIS, J. (2009). Making Sense of Cents: An Examination of Firms That Marginally Miss or Beat Analyst Forecasts. *The Journal of Finance*, 64(5):2361–2388.

Bisbee, J., Larson, J., and Munger, K. (2020). #polisci twitter: A descriptive analysis of how political scientists use twitter in 2019. *Perspectives on Politics*, pages 1–22.

Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654.

Bloom, N., Schankerman, M., and Reenen, J. V. (2013). Identifying technology spillovers and product market rivalry. *Econometrica*, 81(4):1347–1393.

Bobtcheff, C., Bolte, J., and Mariotti, T. (2016). Researcher's dilemma. *The Review of Economic Studies*, 84(3):969—-1014.

Boehm, C. E., Levchenko, A. A., and Pandalai-Nayar, N. (2023). The long and short (run) of trade elasticities. *American Economic Review*, 113(4):861–905.

Bolton, P. and Scharfstein, D. S. (1996). Optimal debt structure and the number of creditors. *Journal of Political Economy*, 104(1):1–25.

Borusyak, K., Hull, P., and Jaravel, X. (2021). Quasi-experimental shift-share research designs. *The Review of Economic Studies*, 89(1):181–213.

Boz, E., Casas, C., Georgiadis, G., Gopinath, G., Mezo, H. L., Mehl, A., and Nguyen, T. (2022). Patterns of invoicing currency in global trade: New evidence. *Journal of International Economics*, 136:103604.

Bradley, D., Jame, R., and Williams, J. (2022). Non-Deal Roadshows, Informed Trading, and Analyst Conflicts of Interest. *The Journal of Finance*, 77(1):265–315.

Brechwald, W. A. and Prinstein, M. J. (2011). Beyond homophily: A decade of advances in understanding peer influence processes. *Journal of Research on Adolescence*, 21(1):166–179.

Brogaard, J., Engelberg, J., and Parsons, C. A. (2014). Networks and productivity: Causal evidence from editor rotations. *Journal of Financial Economics*, 111(1):251–270.

Brogaard, J., Engelberg, J., and Wesep, E. V. (2018). Do economists swing for the fences after tenure? *Journal of Economic Perspectives*, 32(1):179–194.

Brown, L. D. (2001). A Temporal Analysis of Earnings Surprises: Profits versus Losses. *Journal of Accounting Research*, 39(2):221–241.

Bruno, V. and Shin, H. S. (2023). Dollar and exports. *The Review of Financial Studies*, 36(8):2963–2996.

Carhart, M. M. (1997). On Persistence in Mutual Fund Performance. *The Journal of Finance*, 52(1):57–82.

Choi, J., Hackbarth, D., and Zechner, J. (2018). Corporate debt maturity profiles. *Journal of Financial Economics*, 130(3):484–502.

Clement, M. B. and Tse, S. Y. (2005). Financial Analyst Characteristics and Herding Behavior in Forecasting. *The Journal of Finance*, 60(1):307–341.

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

Cohen, L., Frazzini, A., and Malloy, C. (2008). The small world of investing: Board connections and mutual fund returns. *Journal of Political Economy*, 116(5):951–979.

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

Cookson, J. A., Engelberg, J. E., and Mullins, W. (2023). Echo Chambers. *The Review of Financial Studies*, 36(2):450–500.

Cookson, J. A. and Niessner, M. (2020). Why don't we agree? evidence from a social network of investors. *The Journal of Finance*, 75(1):173–228.

Coppola, A., Maggiori, M., Neiman, B., and Schreger, J. (2021). Redrawing the map of global capital flows: The role of cross-border financing and tax havens. *The Quarterly Journal of Economics*, 136(3):1499–1556.

Crawford, S. S., Gray, W. R., and Kern, A. E. (2017). Why do fund managers identify and share profitable ideas? *Journal of Financial and Quantitative Analysis*, 52(5):1903–1926.

Crossref (2021). January 2021 public data file. Database, Crossref.

Cuculiza, C., Antoniou, C., Kumar, A., and Maligkris, A. (2021). Terrorist Attacks, Analyst Sentiment, and Earnings Forecasts. *Management Science*, 67(4):2579–2608.

Dai, M., Giroud, X., Jiang, W., and Wang, N. (2024). A q theory of internal capital markets. *The Journal of Finance*.

Dai, R., Donohue, L., Drechsler, Q. F., and Jiang, W. (2022). Dissemination, publication, and impact of finance research: When novelty meets conventionality. *Review of Finance*, 27(1):79–141.

Dechezleprêtre, A., Einiö, E., Martin, R., Nguyen, K.-T., and Van Reenen, J. (2023). Do tax incentives increase firm innovation? an rd design for r&d, patents, and spillovers. *American Economic Journal: Economic Policy*, 15(4):486–521.

Diamond, D. W. (1993). Seniority and maturity of debt contracts. *Journal of Financial Economics*, 33(3):341–368.

Duflo, E. and Saez, E. (2003). The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment. *The Quarterly journal of economics*, 118(3):815–842.

Edmans, A., García, D., and Norli, O. (2007). Sports Sentiment and Stock Returns. *The Journal of Finance*, 62(4):1967–1998.

Eiteman, D. K., Stonehill, A. I., and Moffett, M. H. (2015). *Multinational Business Finance*. Pearson.

Elton, E. J. and Gruber, M. J. (1972). Earnings Estimates and the Accuracy of Expectational Data. *Management Science*, 18(8):B–409.

Elton, E. J., Gruber, M. J., and Gultekin, M. (1981). Expectations and Share Prices. *Management Science*, 27(9):975–987.

Eren, O. and Mocan, N. (2018). Emotional Judges and Unlucky Juveniles. *American Economic Journal: Applied Economics*, 10(3):171–205.

Fama, E. F. and French, K. R. (2006). Profitability, investment and average returns. *Journal of Financial Economics*, 82(3):491–518.

Fama, E. F. and French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, 116(1):1–22.

Fang, B. and Hope, O.-K. (2020). Analyst teams. *Review of Accounting Studies*, 26(2):425–467.

Filippas, A. and Horton, J. (2021). The production and consumption of social media. Working Paper, National Bureau of Economic Research.

Fischer, L. F. (2022). Sharing is caring? knowledge diffusion in researcher networks. *Working Paper*.

Fischer, L. F. (2024). Of coupons and cargo: International debt, production, and trade. *Working Paper*.

Fischer, L. F. and Shore, E. P. (2023). 'did you catch the game last night?' peer group effects in sell-side analyst forecasts. *Working Paper*.

Fitzgerald, D. and Haller, S. (2018). Exporters and shocks. *Journal of International Economics*, 113:154–171.

Flaaen, A., Haberkorn, F., Lewis, L., Monken, A., Pierce, J., Rhodes, R., and Yi, M. (2023). Bill of lading data in international trade research with an application to the COVID-19 pandemic. *Review of International Economics*, 31(3):1146–1172.

Fourcade, M., Ollion, E., and Algan, Y. (2015). The superiority of economists. *Journal of Economic Perspectives*, 29(1):89–114.

Frankel, R., Kothari, S. P., and Weber, J. (2006). Determinants of the informativeness of analyst research. *Journal of Accounting and Economics*, 41(1):29–54.

Fried, D. and Givoly, D. (1982). Financial analysts' forecasts of earnings: A better surrogate for market expectations. *Journal of Accounting and Economics*, 4(2):85–107.

Friedrich, B. U. and Zator, M. (2023). Flexibility costs of debt: Danish exporters during the cartoon crisis. *Journal of Financial Economics*, 148(2):91–117.

Garicano, L. and Santos, T. (2004). Referrals. *American Economic Review*, 94(3):499–525.

Gibbons, B., Iliev, P., and Kalodimos, J. (2020). Analyst Information Acquisition via EDGAR. *Management Science*, 67(2):769–793.

Gillan, S. L., Koch, A., and Starks, L. T. (2021). Firms and social responsibility: A review of esg and csr research in corporate finance. *Journal of Corporate Finance*, 66:101889.

Givoly, D. and Lakonishok, J. (1979). The information content of financial analysts' forecasts of earnings: Some evidence on semi-strong inefficiency. *Journal of Accounting and Economics*, 1(3):165–185.

Givoly, D. and Lakonishok, J. (1980). Financial analysts' forecasts of earnings: Their value to investors. *Journal of Banking & Finance*, 4(3):221–233.

Gleason, C. A. and Lee, C. M. C. (2003). Analyst Forecast Revisions and Market Price Discovery. *The Accounting Review*, 78(1):193–225.

Gopinath, G., Itskhoki, O., and Rigobon, R. (2010). Currency choice and exchange rate pass-through. *American Economic Review*, 100(1):304–336.

Gozzi, J. C., Levine, R., Martinez Peria, M. S., and Schmukler, S. L. (2015). How firms use corporate bond markets under financial globalization. *Journal of Banking & Finance*, 58:532–551.

Griffin, P. A. (1976). Competitive Information in the Stock Market: An Empirical Study of Earnings, Dividends and Analysts' Forecasts. *The Journal of Finance*, 31(2):631–650.

Hall, B., Jaffe, A., and Trajtenberg, M. (2001). The NBER patent citation data file: Lessons, insights and methodological tools. Working Paper, NBER.

Hall, B. and Reenen, J. V. (2000). How effective are fiscal incentives for r&d? a review of the evidence. *Research Policy*, 29(4-5):449–469.

Hall, B. H. (1995). Effectiveness of research and experimentation tax credits: Critical literature reviewand research design. Technical report, Office of Technology Assessment, Congress of the United States.

Handy, C. (2007). *Understanding Organizations*. Penguin UK.

Harford, J., Jiang, F., Wang, R., and Xie, F. (2019). Analyst Career Concerns, Effort Allocation, and Firms' Information Environment. *The Review of Financial Studies*, 32(6):2179–2224.

Hassan, T. A., Schreger, J., Schwedeler, M., and Tahoun, A. (2023). Sources and transmission of country risk. *Review of Economic Studies*.

Helpman, E., Melitz, M. J., and Yeaple, S. R. (2004). Export versus FDI with heterogeneous firms. *American Economic Review*, 94(1):300–316.

Hill, R. and Stein, C. (2020). Scooped! estimating rewards for priority in science. *Working Paper*.

Hill, R. and Stein, C. (2021). Race to the bottom: Competition and quality in science. *Working Paper*.

Hong, H. and Kubik, J. D. (2003). Analyzing the Analysts: Career Concerns and Biased Earnings Forecasts. *The Journal of Finance*, 58(1):313–351.

Hong, H., Kubik, J. D., and Stein, J. C. (2004). Social interaction and stock-market participation. *The Journal of Finance*, 59(1):137–163.

Hong, H. and Shore, E. P. (2022). Corporate social responsibility. Working Paper, National Bureau of Economic Research.

Horst, T. (1972). The industrial composition of u. s. exports and subsidiary sales to the canadian market. *American Economic Review*, 62(1):37–45.

Imhoff, E. A. and Lobo, G. J. (1984). Information Content of Analysts' Composite Forecast Revisions. *Journal of Accounting Research*, 22(2):541–554.

Keloharju, M. and Niskanen, M. (2001). Why do firms raise foreign currency denominated debt? evidence from finland. *European Financial Management*, 7(4):481–496.

Kempf, E. and Tsoutsoura, M. (2021). Partisan Professionals: Evidence from Credit Rating Analysts. *The Journal of Finance*, 76(6):2805–2856.

Kim, J., Lee, S. J., and Marschke, G. (2006). International knowledge flows: Evidence from an inventor-firm matched data set. Working Paper, National Bureau of Economic Research.

Kim, J. and Marschke, G. (2010). Teams in r&d: Evidence from inventor data. *Working Paper*.

Kim, J.-B., Lu, L. Y., and Yu, Y. (2018). Analyst Coverage and Expected Crash Risk: Evidence from Exogenous Changes in Analyst Coverage. *The Accounting Review*, 94(4):345–364.

Kong, D., Lin, Z., Wang, Y., and Xiang, J. (2021). Natural disasters and analysts' earnings forecasts. *Journal of Corporate Finance*, 66:101860.

Kreikebaum, H. (2003). *Organisationsmanagement internationaler Unternehmen*. Lehrbuch. Gabler, Wiesbaden, 2., vollst. überarb. und erw. aufl., nachdr. edition. Hier auch später erschienene, unveränderte Nachdrucke.

Leland, H. E. (1994). Corporate debt value, bond covenants, and optimal capital structure. *The Journal of Finance*, 49(4):1213–1252.

Li, K., Mai, F., Shen, R., and Yan, X. (2021). Measuring corporate culture using machine learning. *The Review of Financial Studies*, 34(7):3265–3315.

Liao, G. Y. (2020). Credit migration and covered interest rate parity. *Journal of Financial Economics*, 138(2):504–525.

Lim, T. (2001). Rationality and Analysts' Forecast Bias. *The Journal of Finance*, 56(1):369–385.

Liu, X., Tang, H., Wang, Z., and Wei, S.-J. (2022). Currency carry trade by trucks: The curious case of china's massive imports from itself. *Review of Finance*, 27(2):469–493.

Madrian, B. C. and Shea, D. F. (2001). The power of suggestion: Inertia in 401 (k) participation and savings behavior. *The Quarterly journal of economics*, 116(4):1149–1187.

Maggiori, M., Neiman, B., and Schreger, J. (2020). International currencies and capital allocation. *Journal of Political Economy*, 128(6):2019–2066.

Manova, K. (2012). Credit constraints, heterogeneous firms, and international trade. *The Review of Economic Studies*, 80(2):711–744.

Manova, K., Wei, S.-J., and Zhang, Z. (2015). Firm exports and multinational activity under credit constraints. *Review of Economics and Statistics*, 97(3):574–588.

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6):1695–1725.

Modigliani, F. and Miller, M. H. (1958). The cost of capital, corporation finance and the theory of investment. *American Economic Review*, 48(3):261–297.

MSCI (2023). Msci fixed income data methodology. Technical report, MSCI Inc.

Murray, F., Aghion, P., Dewatripont, M., Kolev, J., and Stern, S. (2016). Of mice and academics: Examining the effect of openness on innovation. *American Economic Journal: Economic Policy*, 8(1):212–252.

Müller, R. (2015). A technical note on the svensson modelas applied to the swiss term structure. Technical report, Bank for International Settlements.

Nelson, C. R. and Siegel, A. F. (1987). Parsimonious modeling of yield curves. *The Journal of Business*, 60(4):473.

Park, C. (2000). Monitoring and structure of debt contracts. *The Journal of Finance*, 55(5):2157–2195.

Petre, D. (2009). Mining individual securities databases for analytical purposes: the example of the bis international debt securities statistics. In Settlements, B. f. I., editor, *Proceedings of the IFC Workshop on*, volume 29, pages 177–181. Bank for International Settlements.

Richmond, R. J. (2019). Trade network centrality and currency risk premia. *The Journal of Finance*, 74(3):1315–1361.

Roberts, M. R. and Sufi, A. (2009). Control rights and capital structure: An empirical investigation. *The Journal of Finance*, 64(4):1657–1695.

Shin, H.-H. and Stulz, R. M. (1998). Are internal capital markets efficient? *The Quarterly Journal of Economics*, 113(2):531–552.

Shore, E. (2023). Living up to Expectations: Sell-side Analyst Forecasts and Firm Behavior. *Working Paper*.

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

So, E. C. (2013). A new approach to predicting analyst forecast errors: Do investors overweight analyst forecasts? *Journal of Financial Economics*, 108(3):615–640.

Stein, J. C. (2008). Conversations among competitors. *American Economic Review*, 98(5):2150–2162.

Stock, J. H., Wright, J. H., and Yogo, M. (2002). A survey of weak instruments and weak identification in generalized method of moments. *Journal of Business & Economic Statistics*, 20(4):518–529.

Stock, J. H. and Yogo, M. (2002). Testing for Weak Instruments in Linear IV Regression. Working Paper 284, National Bureau of Economic Research.

Svensson, L. E. O. (1994). Estimating and interpreting forward interest rates: Sweden 1992-1994. *IMF Working Papers*, 94(114):1.

Terry, S. (2015). The Macro Impact of Short-Termism. Discussion Paper 15-022, Stanford Institute for Economic Policy Research.

Trueman, B. (1994). Analyst Forecasts and Herding Behavior. *The Review of Financial Studies*, 7(1):97–124.

van Binsbergen, J. H., Diamond, W. F., and Grotteria, M. (2022). Risk-free interest rates. *Journal of Financial Economics*, 143(1):1–29.

Wang, Z., Ettinger, M., Xie, Y., and Xu, L. (2020). The cost of capital: U.s.-based multinational corporations versus u.s. domestic corporations. *Global Finance Journal*, 44:100443.

Wilson, D. J. (2009). Beggar thy neighbor? the in-state, out-of-state, and aggregate effects of r&d tax credits. *Review of Economics and Statistics*, 91(2):431–436.

Appendix A: Of Coupons and Cargo

A.1 S&P Panjiva

Table A.1: S&P Panjiva – Overview

This table provides an overview of the coverage in S&P Panjiva. The exact variables available depend on the reporting country. The data is cleaned in a way that a common set of variables, which are available for the majority of countries in the data, is kept.

Country	Modes	History	HS	$ Values	Notes
Bolivia	All	2014 - Aug 2018	11	Y	
Brazil	Sea	2014 - Present	8	N	
Chile	All	2009 - Present	8	Y	
China	All	2013 - 2018	8	Y	Not transaction level
Colombia	All	2007 - Present	6	Y	
Costa Rica	All	2014 - Present	10	N	
Ecuador	Sea	2014 - Present	12	N	
India	All	2016 - Present	-	Y	
Sri Lanka	All	2016 - Present	8	Y	
Mexico	All	2011 - Present	8	Y	
Pakistan	All	2018 - Present	10	Y	
Panama	All	2009 - Present	8	Y	
Paraguay	All	2014 - Present	8	Y	
Peru	All	2011 - Present	10	Y/N	
US	Sea	2009 - Present	6	Estimate	Imports from 2007
Uruguay	All	2003 - Present	10	Y	
Venezuela	All	2014 - 2019	10	Y	

Figure A.1: S&P Panjiva Coverage by Country

This figure depicts the years for which trade data is available for all of the 16 countries. The bars are greyed out if the data for the corresponding period is only available for less than an entire year and/or does not cover imports as well as exports. China stopped making import/export data available at the onset of the China-U.S. trade war in late 2017.

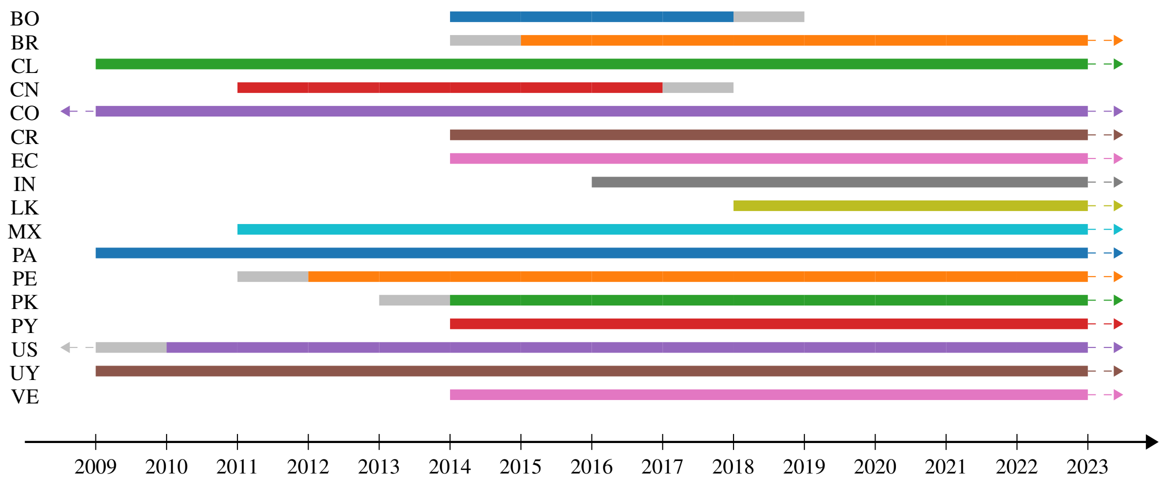

Figure A.2: S&P Panjiva vs. UN Comtrade

This figure presents coefficient estimates and confidence intervals from the following regression $\ln(trade_{c,t}^{S\&P}) = \beta \ln(trade_{c,t}^{UN}) + \lambda_t + \varepsilon_t$. The dependent variable is the log of the total trade reported between the focal country and trading partner c in S&P Panjiva, the independent variable is the corresponding number taken from UN Comtrade. Error bars represent 95% confidence intervals around the point estimated, based on White standard errors. The disaggregated, shipment-level quantities recover more than 75% of total trade. Differences arise due to the coverage in S&P Panjiva being limited to containerized and/or maritime freight only in the case of certains countries (see Table A.1).

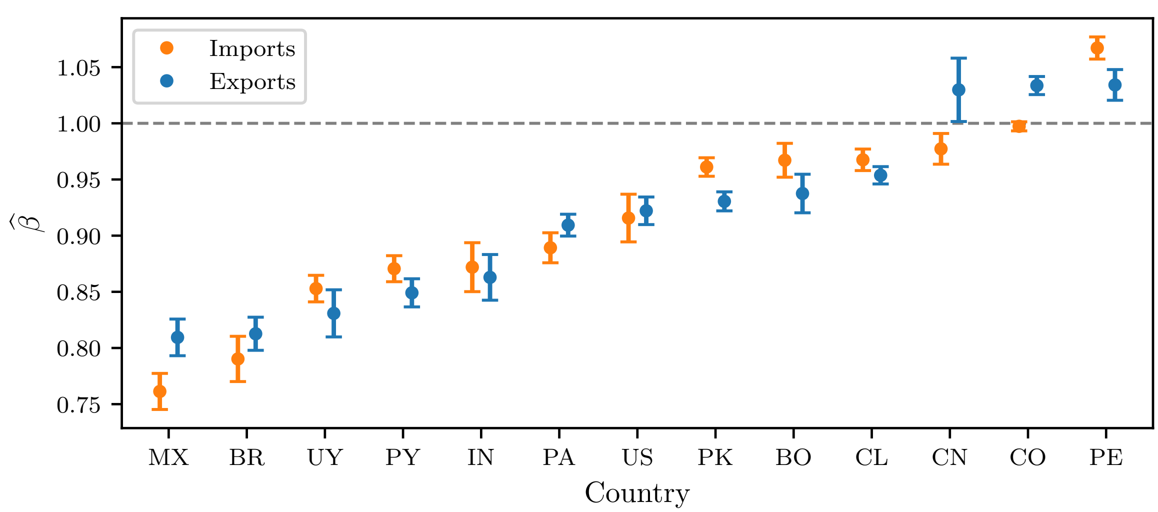

A.2 Refinitiv

Table A.2: Refinitiv-Compustat-Merged

This figure presents the coverage of the merged Refinitiv-Compustat data. All instruments, regardless of their issuer, are included in the column labelled 'Raw'. Only issues that could be linked to the ultimate issuer in Compustat are included in the column labelled 'Linked'. The final column, labelled 'Coverage' is computed as the share of all issues and face value that could be linked to an entity in Compustat.

Location	**Raw**		**Linked**		**Coverage**	
	#Issues	bn USD	#Issues	bn USD	#Issues	bn USD
EU	2,129,162	40.44	1,157,549	27.34	0.54	0.68
CN	194,425	18.03	55,693	13.41	0.29	0.74
US	1,425,676	12.01	730,558	8.93	0.51	0.74
KR	228,280	2.48	160,029	2.11	0.70	0.85
AU	46,694	1.56	17,534	0.71	0.38	0.45
JP	12,008	1.51	7,212	1.31	0.60	0.87
IN	61,941	0.73	34,478	0.57	0.56	0.78
CA	195,405	0.63	81,831	0.43	0.42	0.68
HK	12,350	0.62	3,231	0.29	0.26	0.47
RU	4,243	0.47	1,554	0.36	0.37	0.77
MX	19,424	0.44	8,244	0.36	0.42	0.82
Other	675,354	3.61	210,691	2.56	0.31	0.71

Table A.3: Corporate Events

This table reports the different corporate events leading to changes in the amount outstanding of the fixed income securities in the dataset. The change in amount is relative to the nominal value at issuance. For a detailed description of the events see MSCI (2023). At issuance the value therefore is mechanically equal to 1, however, at maturity it may be different from -1 due to funding events having taken place between issuance and maturity.

Event	**N**	**Mean**	**S.D.**	**p25**	**Median**	**p75**
Issued	6,445,377	1.00	0.00	1.00	1.00	1.00
Matured	5,310,267	-0.99	0.09	-1.00	-1.00	-1.00
Called	908,482	-0.98	0.12	-1.00	-1.00	-1.00
Reopened	185,693	0.34	0.31	0.08	0.25	0.53
Repayment before maturity	168,615	-0.11	0.21	-0.10	-0.03	-0.01
Repurchased	164,067	-0.36	0.36	-0.67	-0.23	-0.04
Prepayment	160,012	0.02	0.35	-0.08	-0.00	0.11
Excahnged	62,998	-0.30	0.42	-0.86	-0.02	-0.00
Defaulted	21,184	-0.03	0.09	-0.03	-0.01	-0.00
Funged	8,253	-1.00	0.04	-1.00	-1.00	-1.00
Put	5,780	-0.58	0.41	-1.00	-0.70	-0.13
Tendered	5,484	-0.38	0.29	-0.58	-0.33	-0.13
Not Active	5,001	-0.99	0.11	-1.00	-1.00	-1.00
Increased due to Funding	4,272	0.42	0.27	0.20	0.38	0.60
Liquidated	2,107	-0.90	0.26	-1.00	-1.00	-1.00
Cancelled	1,381	-0.91	0.25	-1.00	-1.00	-1.00
Overallotment	1,129	0.14	0.09	0.10	0.15	0.15
Prepayment - Pro Rata	1,056	-0.67	0.38	-1.00	-1.00	-0.33
Call - Pro Rata	520	-0.19	0.21	-0.29	-0.10	-0.05
Capitalization of Interest	327	0.03	0.04	0.00	0.02	0.04
Restructured	262	-0.72	0.39	-1.00	-1.00	-0.41
Multiple Actions	253	-0.25	0.31	-0.42	-0.15	-0.04
Issued in Exchange	151	0.38	0.30	0.12	0.30	0.64
Retained Amount	55	-0.28	0.17	-0.40	-0.25	-0.20
Write Down	51	-0.91	0.22	-1.00	-1.00	-1.00
Issuance due to a Switch Auction	48	-0.39	0.44	-0.64	-0.42	-0.04
Currency Redenomination	34	-0.80	0.40	-1.00	-0.97	-0.85
Other	16	-1.00	0.00	-1.00	-1.00	-1.00
Remarketed	12	-0.41	0.31	-0.62	-0.50	-0.31
Reverse Auction	7	-0.13	0.12	-0.23	-0.08	-0.03
Refinancing Transaction	5	0.08	0.40	-0.30	0.11	0.26
Write Down - Pro Rata	2	-0.78	0.06	-0.79	-0.78	-0.76

Figure A.3: Share of Discount Instruments

This figure presents the share of instruments that are discount notes weighted by amount issued for different maturity buckets. An security is classified as being a discount instrument if it does not pay any dividends.

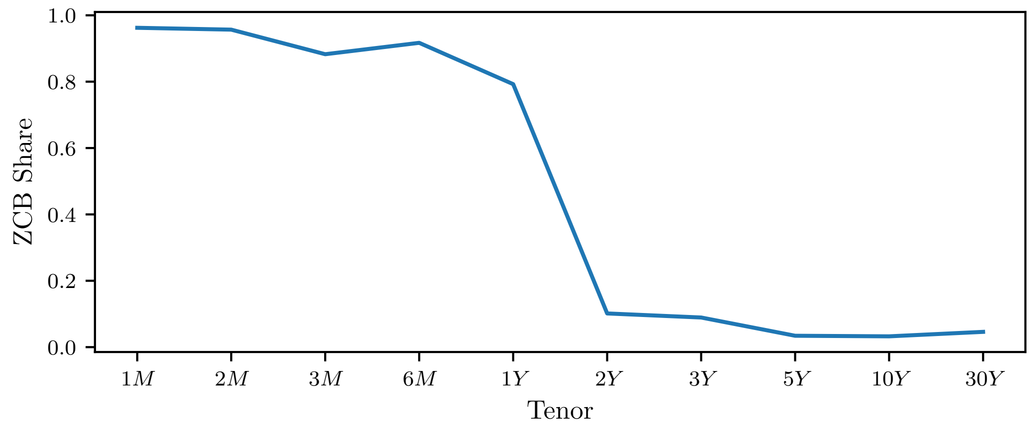

Table A.4: Refinitiv vs. FISD

The table reports coefficient estimates obtained from regressing the most important bond observables from Mergent FISD on the corresponding information as reported in the Refinitiv bond data. Robust standard errors are reported in parentheses. For bonds that are present in the Refinitiv data as well as in Mergent's Fixed Income Securities Database (FISD), the most frequently used bond characteristics in FISD are regressed on their counterpart in Refinitiv. There are no meaningful, systematic discrepancies between the two.

Dependent Variables:	$\ln(amt)$	P	C	ytm_b	$ytm_b - r_f$
Refinitiv	0.9993***	0.9532***	0.9918***	0.9988***	0.9885***
	(2.81×10^{-5})	(0.0019)	(0.0003)	(0.0002)	(0.0014)
Fit statistics					
Observations	282,009	282,176	205,325	77,646	20,956
Adjusted R^2	0.99039	0.94084	0.95245	0.88745	0.93812

Heteroskedasticity-robust standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

A.3 Bureau van Dijk Orbis

Table A.5: Entity Types

This table reports the distribution of entity types as reported in Orbis. The sample consists of all subsidiaries of firms that are in Compustat and which are part of the merged dataset used in this paper. Only those establishments that were active at some point in time between 2010 and 2020 are considered. Subsidiaries classified as Branch or Industrial Company are considered ‘corporate subsidiaries’.

Entity Type	**N**	**Share**	**Cum. Share**
Branch	2,912,912	0.60	0.60
Industrial company	1,675,471	0.35	0.95
Financial company	108,700	0.02	0.97
Fund	75,640	0.02	0.98
Marine vessel	26,825	0.01	0.99
Insurance company	24,560	0.01	0.99
Bank	21,712	0.00	1.00
Foundation/Research institute	4,626	0.00	1.00
Other	996	0.00	1.00
Private equity	928	0.00	1.00
Public authority	97	0.00	1.00
Hedge fund	28	0.00	1.00

A.4 Yield to Maturity

A fixed income security's yield to maturity is given by y, which is the internal rate of return at which the present value of coupon payments and the repayment of the face value coincide with the price at issuance, as outlined below.

$$P = \sum_t \frac{C}{(1+y)^t} + \frac{F}{(1+y)^T}$$

Where P is the price of the security at issuance, C are coupon payments, and F is the face value, which is repaid at maturity. The coupon frequency N defines how many coupon payments are made each year. In the United States, coupon are paid semi-annually, i.e., $N = 2$, whereas European bonds tend to pay interest annually, i.e., $N = 1$. Given coupon rate r_C, the N annual coupon payments are given by C_n as defined below.

$$C_n = \frac{r_C \times F}{N}$$

For tractability, it is helpful to rescale t, such that all periods are integer valued. While not necessary, this trick makes the problem more intuitive and leads to a, minor, increase in performance of the root-finding algorithm. Instead of y, the discount rate solving the expression below, y_n, is the yield for periods of length $\frac{1}{N}$ years.

$$P = \sum_{n=1}^{N \times T} \frac{C_n}{(1+y_n)^n} + \frac{F}{(1+y_n)^{N \times T}}$$

The resulting yield to maturity y_n needs to be compounded N times to arrive at the effective annual YTM.

$$y = (1+y_n)^N - 1$$

Sidenote. In the United States, bonds pay coupons semi-annually and yields are annualized by multiplying by a factor of two. Generalizing this concept, one can alternatively convert y_n as below:

$$y^{US} = N \times y_n$$

Naturally, for small N it holds that $y \approx y^{US}$. Most commercial databases, such as Mergent FISD, report offering yields using this convention. While not problematic when only considering bonds issued in the U.S., in a cross-country analysis is is preferable to use the, mathematically correct, compounding instead.

A.4.1 Zero-Coupon Bonds

Not all types of securities require numerically solving for the origins of a polynomial. Specifically, there are two cases in which there exists a closed form solution to the yield to maturity. The simplest case are securities with $N = 0$, i.e., zero-coupon bonds.

$$y^{ZCB} = \left(\frac{F}{P}\right)^{\frac{1}{T}} - 1$$

A.4.2 Par Bonds

The second special case are securities that are issued *at par*, i.e., where $P = F$. For these instruments the yield to maturity y coincides with the coupon rate r_C. In securities pay interest sub-annually, the yield of a par bond is computed as described below.

$$y^{par} = \left(1 + \frac{r_C}{f}\right)^{f} - 1$$

Notice that the expression above collapses to $y^{par} = r_C$ if coupons are paid annually ($f = 1$). Furthermore, $y^{par} = r_C$ holds for par bonds when applying the U.S. convention of annualizing.

A.4.3 All Other Cases

For all securities that are neither discount notes/zero-coupon bonds, nor issued at par, the yield y is computed by numerically solving the following polynomial for y_n.

$$\sum_{n=1}^{N\times T} \frac{C_n}{(1+y_n)^n} + \frac{F}{(1+y_n)^{N\times T}} - P = 0$$

Followed by compounding y_n to arrive at y. The optimizer in `scipy` is set to Newton's method and provided a starting value of 5%.

A.5 Yield Curve Interpolation

Yield curves are interpolated following Nelson and Siegel (1987) and Svensson (1994). The instantaneous forward rate at time t for m periods is modeled as:

$$f(t, t+m; \boldsymbol{\beta}, \boldsymbol{\tau}) = \beta_0 + \beta_1 \exp\left(-\frac{m}{\tau_1}\right) + \beta_2 \frac{m}{\tau_1} \exp\left(-\frac{m}{\tau_1}\right) + \beta_3 \frac{m}{\tau_2} \exp\left(-\frac{m}{\tau_2}\right)$$

The parameters $\boldsymbol{\beta} = (\beta_0, \beta_1, \beta_2, \beta_3)^\top$ and $\boldsymbol{\tau} = (\tau_1, \tau_2)^\top$ are estimated by minimizing the squared error. Equipped with the estimates, the corresponding yields are derived by integrating over $f(t, t+m)$:

$$\begin{aligned} i(t, t+m; \widehat{\boldsymbol{\beta}}, \widehat{\boldsymbol{\tau}}) =& \widehat{\beta}_0 + \widehat{\beta}_1 \frac{1-\exp\left(-\frac{m}{\widehat{\tau}_1}\right)}{\frac{m}{\widehat{\tau}_1}} + \widehat{\beta}_2 \left(\frac{1-\exp\left(-\frac{m}{\widehat{\tau}_1}\right)}{\frac{m}{\widehat{\tau}_1}} - \exp\left(-\frac{m}{\widehat{\tau}_1}\right)\right) \\ &+ \widehat{\beta}_3 \left(\frac{1-\exp\left(-\frac{m}{\widehat{\tau}_2}\right)}{\frac{m}{\widehat{\tau}_2}} - \exp\left(-\frac{m}{\widehat{\tau}_2}\right)\right) \end{aligned}$$

For a more detailed explanation of the model and how it is estimated, see Müller (2015).

Appendix B: Did You Catch the Game Last Night?

Table B.1: Data Sources

This table lists the data sources, how they were accessed, as well as the month the data was obtained.

Source	**Access Mode**	**Timestamp**
Bloomberg	Bloomberg Terminal (Anywhere)	June 2022
CapitalIQ	CapitalIQ Pro website	September 2022
Compustat	WRDS	January 2023
CRSP	WRDS	January 2023
IBES	WRDS	January 2023
LinkedIn	LinkedIn Website	June 2022
WRDS Link	WRDS	January 2023

Figure B.1: Distribution of Analysts' Alma Maters

The figures below present the distribution of the higher education institution attended by the sell-side analysts in our sample. For each university we count the number of analysts that attended it, inversely weighted by the number of different institutions the analysts went to. We then aggregate to the state level (for US universities, Panel b) and the country level (Panel a).

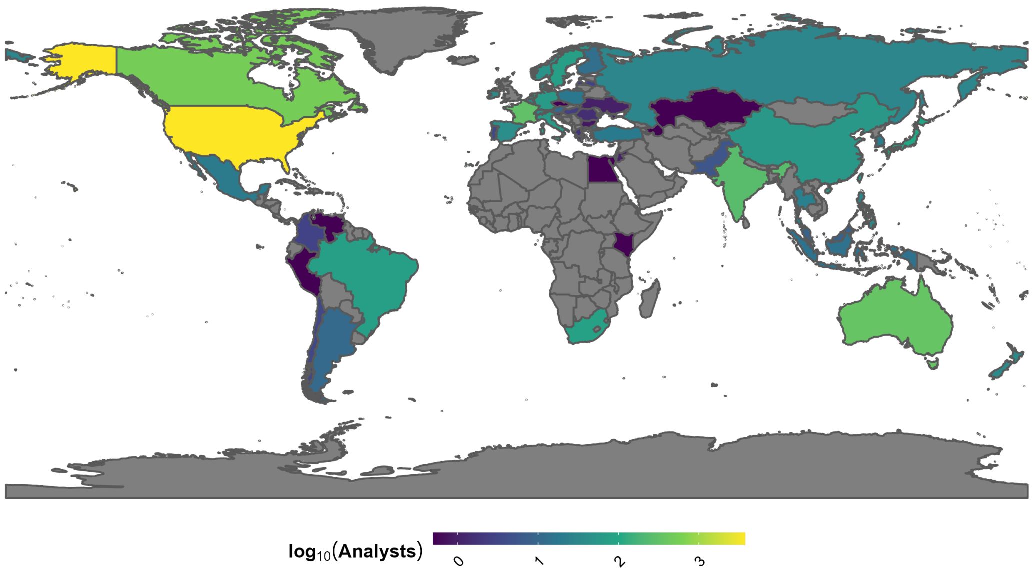

(a) World

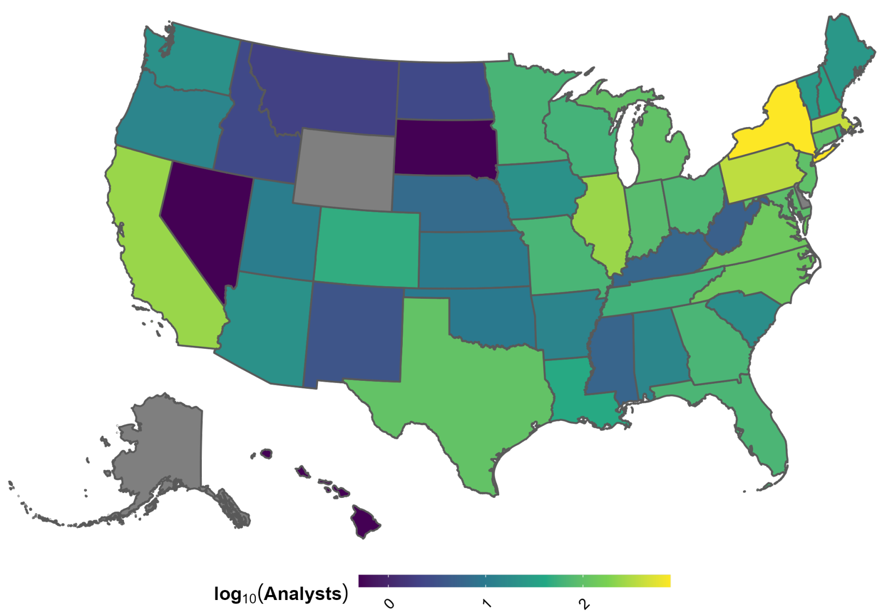

(b) United States

Appendix C: Sharing is Caring?

C.1 JEL Classification

One common way of delineating different fields of research in economics is the classification scheme developed by the Journal of Economic Literature. The current set of 20 areas is presented in Table C.1 below. These definitions are used to assign individual papers to research areas, as well as across publications to classify the authors based on their work.

Table C.1: JEL Classification System

Classification system developed for the *Journal of Economic Literature* (JEL). Each subfield (or group of subfields) is assigned an alphabetical code. The list of codes and corresponding fields, which is available on the website of the *American Economic Association* (AEA), is reproduced below. Some of the descriptions have been shortened for ease of legibility. The full list, including all sub-classifications, is available at `https://www.aeaweb.org/econlit/jelCodes.php`.

Code	Description
A	General Economics and Teaching
B	History of Economic Thought, Methodology, and Heterodox Approaches
C	Mathematical and Quantitative Methods
D	Microeconomics
E	Macroeconomics and Monetary Economics
F	International Economics
G	Financial Economics
H	Public Economics
I	Health, Education, and Welfare
J	Labor and Demographic Economics
K	Law and Economics
L	Industrial Organization
M	Business Administration/Business Economics, Marketing, Accounting
N	Economic History
O	Economic Development, Innovation, Technological Change, and Growth
P	Political Economy and Comparative Economic Systems
Q	Agricultural/Natural Resource, Environmental/Ecological Economics
R	Urban, Rural, Regional, Real Estate, and Transportation Economics
Y	Miscellaneous Categories
Z	Other Special Topics

C.2 SSRN

Data from the Social Sciences Research Network (SSRN) is obtained using the site's undocumented API. This appendix aims to provide a high-level summary of the API and how it can be applied to obtain details on the working papers uploaded to a given SSRN network. The api is available at the endpoint

`https://api.ssrn.com/content/v1/bindings/:network_no/papers`

where `:network_no` must be replaced with the appropriate network number. Relevant for this paper are the Economics (205) and Financial Economics (203) Networks. A simple request to the API without any further parameters will return the 50 newest papers in the network.

To obtain the complete history of filings the following parameters can be handed over to the API in the usual format:[1] (1) `sort`, (2) `index`, and (3) `count`.

In order to explain these parameters it is important to develop an understanding of how the API functions first. When called, the server behind the API constructs a table with all information available on the papers in network `network_no`. This table is sorted depending on how `sort` is set (0: new to old, 1: old to new, 2: downloads descending, 3: downloads ascending, 4: Z-A, 5: A-Z).

The remaining parameters, `index`, and `count`, the determine which slice of the bigger table is returned to the user. `count` determines the number of results to be returned and can take values between 1 and 200. `index` refers to the row number of the first entry to be returned. In order to obtain the entire list of papers it is thus necessary to send repeat requests and concatenate the result. As is common for APIs there are access limits which as of now are unknown to me. Implementing a pause of 5 seconds between API calls lead to most calls being successful.

The API has one major shortcoming as it can only deal with `index` of values up to $400,000$. To get the total number of papers in a given network, the header of each valid API call contains

[1] See https://www.dataquest.io/blog/r-api-tutorial/ for an introduction on how to use APIs for data access in R and https://www.dataquest.io/blog/python-api-tutorial/ for a the same in Python.

the number of papers (total). Thus, if a given network contains more than 400, 000 but less than 800, 000 papers one has to access the list in ascending and descending order and then combine the results.

The API returns the following variables: id, abstract_type, publication_status, is_paid, reference, page_count, title, authors, affiliations, is_approved, approved_date, url, and downloads.

C.3 NBER

Data from the NBER is retrieved starting from the list of working papers using Selenium web scraping. For each paper the authors, the classification of the paper, the upload date, and the link to the paper are recorded. Based on the list of authors their respective profiles are scraped. For each author their name, affiliation, and position at the NBER are stored. There are several fields which are relevant to the construction of the sample that are optional and thus not available for all individuals. These include the list to the author's homepage, their email address, their phone number, their Twitter handle, and the link to their profile on RePEc.

C.4 RePEc

RePEc (Research Papers in Economics) maintains a database of economists and their work. While broadly similar to SSRN, the project is free and open, whereas SSRN is a commercial product. Researchers generally face a trade-off when choosing how to make their ongoing research available. In terms of the relevance for this project, the major difference between the two is that while RePEc does track journal publications as well as working papers, SSRN is limited to research at the working paper stage. Data from RePEc is downloaded into a database using a modified version of the python package RePEc written by Andrei Dubovik.[2]

Publication records are then linked to authors using the RePEc person data.[3] The person ref

[2]Available at https://github.com/andrei-dubovik/RePEc. Modifications include the extraction of the optional doi field from the ReDIF files.

[3]Available at ftp://all.repec.org/pub/RePEc/per/pers/

files contain the metadata of each individual as well as a list of all working and published papers they (co-)authored.

C.5 Crossref

Citation records are obtained from Crossref (2021). The paper metadata and citations are current as of January 7^{th} 2021.[4] As Crossref constructs the citation records from the `doi` of the publications, this variable represents the unique identifier in this dataset.

[4]https://academictorrents.com/details/e4287cb7619999709f6e9db5c359dda17e93d515