NBER WORKING PAPER SERIES

THE VALUE OF INTERMEDIATION IN THE STOCK MARKET

Marco Di Maggio
Mark L. Egan
Francesco Franzoni

Working Paper 26147
http://www.nber.org/papers/w26147

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2019, Revised July 2021

We thank Malcolm Baker, Vyacheslav Fos, Robin Greenwood, Boris Groysberg, Paul Irvine, Vincent van Kervel, Brian Levine, Gregor Matvos, Amit Seru, Daniel Strack, Vish Viswanathan and Mao Ye for their comments and the seminar participants at the Boston Conference on Markets and Competition, the Columbia Workshop in New Empirical Finance, Erasmus University, the Finance UC International Conference at Catholic University of Chile, Harvard Business School, Maastricht University, Norges Bank, The Stanford Institute for Theoretical Economics Financial Regulation Workshop, University California Berkeley Haas School of Business, UNC-Duke Corporate Finance Conference, and at the Western Finance Association Meetings. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

The Value of Intermediation in the Stock Market
Marco Di Maggio, Mark L. Egan, and Francesco Franzoni
NBER Working Paper No. 26147
August 2019, Revised July 2021
JEL No. G12, G23, G24, G28

ABSTRACT

We estimate a structural model of broker choice to quantitatively decompose the value that institutional investors attach to broker services. Studying over 300 million institutional equity trades, we find that investors are sensitive to both explicit and implicit trading costs and are willing to pay a premium for access to formal and informal research. Formal and informal research account for roughly half of the value generated by brokers. Lastly, we use our model to investigate soft-dollar arrangements, where research and execution services are bundled, and find that such arrangements allow hedge funds and mutual funds to underreport management fees by 10%.

Marco Di Maggio
Harvard Business School
Baker Library 265
Soldiers Field
Boston, MA 02163
and NBER
mdimaggio@hbs.edu

Mark L. Egan
Harvard Business School
Baker Library 365
Soldiers Field
Boston, MA 02163
and NBER
megan@hbs.edu

Francesco Franzoni
Swiss Finance Institute
Via G. Buffi 13
6904, Lugano - Switzerland
and University of Lugano
francesco.franzoni@usi.ch

1. Introduction

Most institutional investors do not access equity markets directly. Rather, the majority of institutional investors rely on "high-touch" (non-electronic) brokers, where trading orders are often placed over the phone.[1] Given the development of, in principle, cheaper trading alternatives, such as direct market access and alternative trading systems, why do institutional investors continue to execute trades through brokers? Brokers offer a variety of services to investors and create value by providing efficient execution, market research, and order flow information. Traditionally, brokers have bundled these services into one package, and asset managers have paid for these services with soft dollars (i.e., through trading commissions) rather than paying for them directly (i.e., hard dollars), which makes it difficult to disentangle how brokers create value for investors.

The use of soft dollars also raises transparency issues about the cost of asset management. The concern with soft-dollar commission payments is that they are borne by the end-investor and are opaque, as the value of such soft-dollar payments is not explicitly reported and itemized by the fund unlike other fund expenses, such as management and distribution fees.[2] These transparency issues have attracted the attention of regulators and policymakers and have subsequently resulted in several recent policy interventions, such as The Markets in Financial Instruments Directive (MiFID) II, which aim to hold investment managers accountable to best execution standards and offer greater transparency about the services brokers provide to investors and the associated expenses.[3] In general, the lack of transparency makes it hard to assess whether investment managers abide by their mandate of best execution.[4]

In this paper, we develop and estimate a quantitative framework for understanding how institutional investors execute equity trades using rich microdata. Our empirical framework allows us to address two important questions in the literature that are central to the policy debate. First, we quantify the value that institutional investors give to the different services offered by brokers. For example, a long literature documents that sell-side research analysts impact trading decisions, but how much do investors value research relative to other services offered by the brokers? The answer to this question provides information on how institutions interpret their mandate for best

[1] [https://www.greenwich.com/equities/voice-trading] accessed 5/9/2019

[2] Reporting requirements, such as the SEC Form ADV Part 2, require that funds report whether they pay for services with "soft dollars." Funds must describe the nature of the services they received and the potential of conflicts of interest. However, firms are not required to report the magnitude of soft-dollar transactions (e.g. https://www.sec.gov/about/forms/formadv-part2.pdf). Some asset management firms report the amount of trading commissions that pay for research services in the Statement of Additional Information, a supplement to the Prospectus. Reporting this information is not, however, mandatory.

[3] The Markets In Financial Instruments Directive (MiFID) II was rolled out on January 3, 2018. It applies to all European asset managers, but it has repercussions on global brokerage firms and investment managers selling services to European clients. Before MiFID II, research costs were often 'bundled' into opaque transaction fees borne by funds' clients. Investment managers are now required to pay for research separately from execution services, and either charge clients transparently or pay for research themselves.

[4] According to the SEC the determinative factor of best execution is not just the lowest possible commission cost, but requires that "a money manager should consider the full range and quality of a Broker's services in placing brokerage including, among other things, the value of research provided as well as execution capability, commission rate, financial responsibility, and responsiveness to the money manager." https://www.sec.gov/rules/interp/34-23170.pdf

execution. Moreover, our rich data set and framework allow us to quantify the value of broker services that are less well studied in the literature, such as the value of relationships. Investors are likely to form long-term relationships with their brokers, which affect their sensitivity to fees and propensity to trade with other intermediaries (Goldstein, Irvine, Kandel, and Wiener, 2009). Second, we use our framework to quantify the magnitude of soft dollars in the industry and determine how much higher management fees would be if investment managers paid for broker services with hard, rather than soft, dollars. These results can inform the debate around the recent proposals to unbundle payments in the brokerage industry.

A key challenge to studying these issues has been the lack of data since the analysis requires detailed information on brokerage firms and institutional trading patterns. We can overcome this challenge thanks to a rich micro-data set covering hundreds of millions of equity transactions. Our primary data set comes from Abel Noser Solutions, formerly Ancerno Ltd. The company performs transaction cost analysis for institutional investors and makes the data available for academic research under the agreement of non-disclosure of institutional identity. Our sample covers the period between 2001 and 2014 and includes trade-level data for institutional investors, accounting for up to 20% of the institutional trading volume in the U.S. stock market (Puckett and Yan (2011), Hu, Jo, Wang, and Xie (2018)). Importantly, we observe the identity of the investment manager placing the trade and the broker executing the corresponding trade. We merge the Ancerno sample with brokerage firm-level data from several sources. From these sources, we extract broker characteristics that allow us to proxy for the different dimensions of the quality of services that brokers offer to their clients.[5]

To understand how institutional investors make execution decisions, we develop an empirical model of brokerage firm choice. We examine an investor's trading decision process with a particular emphasis on *where* investors decide to execute their trades. We model an investor's execution decision as a discrete choice problem. Investors choose the broker that maximizes their expected trading profits. When deciding among brokers, investors trade off explicit trading costs (i.e., commissions/fees), implicit trading costs (i.e., price impact), and the quality of other services provided by the broker such as research and order flow information. The framework allows us to measure each investor's subjective valuation of the different services brokers provide.

We estimate this discrete choice framework following the workhorse models used in the industrial organization literature (Berry (1994) and Berry et al. (1995)). Our setting and data are ideal for such estimation for several reasons. First, we observe individual institutional investors making tens of thousands of execution decisions. This rich data allows us to estimate our discrete choice model at the investor level without imposing any parametric assumptions over the distribution of

[5]Note that we do not directly observe the amount of a specific service that a given broker provides to a given client in a given trade, which is part of the motivation for our analysis. For example, we do not know how much research on Apple investment manager A acquired from broker B before trading its stock or the value that manager A placed on that research. Rather, we know how many analysts cover the technology sector at broker B, and whether they are the top-rated analysts in the sector. We then use a revealed preference approach to measure how these broker characteristics, such as the scope and quality of the research coverage at broker B, influence the execution decisions of investment managers, which allows us to determine each investment manager's subjective valuation of broker-provided services.

investor preferences. Second, a common problem in the demand estimation literature is the endogeneity of prices, or in this case trading fees. If brokerage firms adjust their fees in response to the actions and preferences of investors, fees will be endogenous. We address the potential endogeneity of fees through an instrumental variables approach that exploits unique institutional features of the brokerage industry. Brokerage firms charge fees in terms of cents per share traded (Goldstein et al. (2009)). This rigidity in the way fees are set provides exogenous variation in the effective transaction costs paid by investors. To be clear, our identification approach relies on the fact that fees are not expressed as a fraction of the stock price, but rather as cents per share. This practice generates exogenous variation in the price of the transaction.

To summarize, the estimation of our model allows us to quantify the different dimensions of broker-provided services. For the average trade in our sample, formal research accounts for 30% of the value created by brokers, informal research (notably, order-flow information) accounts for 20% of the value created by brokers, and trade execution accounts for the remaining 50% of the value created by brokers (see Figure 3a). Thus, informal services are almost as valuable as formalized research, and, collectively, the different dimensions of research that brokers provide create similar value for asset managers as trade execution. Finally, we exploit the time-series dimension of the data to highlight whether new trends have emerged in recent years.

In more detail, we first examine the price sensitivity of investors. The average broker fee in our data is roughly 3 cents per share or roughly 13bps relative to the value of the transaction. Our broker choice estimates suggest that the majority of institutional investors are relatively price insensitive. The average demand elasticity in our data set is roughly 0.47, which implies that if a broker increases the fee it charges by 1%, its trading volumes will go down by an associated 0.47%. The estimates suggest that investor-broker relationships are "sticky" and that there are many other non-price factors that influence broker choice. However, we find evidence that investors have become more price sensitive in recent years; the elasticity of demand doubled over our sample period (2001 to 2014) as investment managers became more sensitive to trading costs. This finding is consistent with increasing competition in the brokerage industry arising from electronic trading.

In addition to paying explicit execution fees, brokers also face implicit trading costs, which may play an equally important role as explicit fees in driving execution decisions (Anand et al. (2012)). Brokers may differ in their ability to execute large trade orders without moving the market price of a stock. We measure these implicit trading costs at the trade level as the price impact of the trade, i.e. the execution price relative to the price of the stock at the placement of the investor's order. We find that a one standard deviation decrease in price impact is worth 6bps, which is equal to roughly one-half of a standard deviation in broker fees. This finding suggests that both explicit and implicit trading costs play an important role in execution decisions.

Brokers, in addition to execution, offer formal research to their clients, employing equity analysts who provide forecasts, research reports, and general expertise in a given sector. We test whether investors value this broker-provided sell-side research when executing trades. Our estimates indicate that investors are willing to pay a 10-15% higher broker fee (1-2bps relative to the value of the transaction) to have access to a research analyst and an additional 20-40% (3-5bps relative to the value of the transaction) to have access to a top-rated analyst in the sector (as rated by Institutional Investor).

A recent emerging view in the academic literature is that brokers also play the role of information hubs (e.g., Boyarchenko et al. (2021)) because they are likely to have a more comprehensive view of market trends and investors' strategies, and institutional investors likely value this information in all markets. We measure order flow information, which we label informal research, in two ways. Following Di Maggio, Franzoni, Kermani, and Sommavilla (2019), we first define a broker as being informed if he has traded with an informed investor. Second, in line with these authors, we capture the broker's access to information with its centrality in the network of relationships between managers and brokers. We find that investors value these types of informal information as much as formal research. Collectively, we find that formal and informal research account for roughly half of the value that brokers create for investors.

Using equity trader-level information from BrokerCheck, a website operated by the Financial Industry Regulatory Authority (see Egan et al. (2019)), we also have data on the individual traders employed by the corresponding brokerage firm. We find that investors are less likely to trade through a brokerage firm whose equity traders are involved in more client disputes and regulatory offenses and are more likely to trade through a brokerage firm with more experienced traders.[6] Lastly, we find evidence that investors prefer to trade through equity traders located in the same city as the investor. Even though the equity orders are placed either electronically or over the phone, physical proximity to the broker influences an investor's trading decision. This is consistent with the idea that "trading is—and always has been—a relationship business."[7] Moreover, this finding extends the evidence of local bias in asset management (Coval and Moskowitz (1999)) to the choice of trading intermediaries.

This rich setting also allows us to explore how the execution decisions and preferences vary across investors. We do not expect all institutional investors to give the same value on broker-provided services. For example, hedge funds and index funds may not value sell-side research to the same degree as large active mutual funds. Consistent with this intuition, while we find that the average investor values sell-side equity research, we also find that roughly one-third of investors place no value on sell-side research. Hedge funds, as opposed to mutual funds, place a lower value on sell-side research. Conversely, hedge funds appear to place a premium on informed order flow and prefer to trade with brokers that are less centrally located within the trading network, consistent with Ye and Zhu (2019). Similarly, we find that, as expected, index funds do not choose their brokers based on research, which is also useful in validating our empirical framework by making it less plausible for unobserved characteristics of the brokers to be driving our results.

[6] Roughly 6.5% of the traders in our sample have a record of misconduct, which includes customer disputes resulting in a settlement and regulatory offenses. Following Egan et al. (2019) we define misconduct as any customer dispute that resulted in a settlement, regulatory offenses, criminal offenses, and cases where the trader was fired for cause.

[7] The quote is from Johnson, Vice President of Market Structure and Technology at Greenwich Associates. [https://www.bloomberg.com/professional/blog/human-high-touch-trading-stay/] accessed 5/9/2019.

While brokerage firms have traditionally bundled their services, the industry has slowly moved away from bundling over the last fifteen years. As part of recent changes in regulations corresponding to MiFID II, European regulators mandate that brokers must unbundle their services. Previous research has shown that bundled commissions are associated with greater agency costs and have more adverse return consequences than more transparent fund expenses (Edelen et al. (2012)). The concern with soft-dollar payments is that they are borne by the investor but are not explicitly disclosed by the fund; however, it is usually challenging to quantify this type of underreporting. We overcome this challenge by separately calculating the investment manager's shadow-value of broker-produced sell-side research following the methodology used in Petrin (2002).[8] Our estimates suggest that if investment managers had to pay for research in terms of hard rather than soft dollars, investment management fees would be up to 10% higher. There is also substantial heterogeneity across managers; our estimates suggest that the use of soft dollars potentially allows 25% of institutional investors, weighted by assets, to underreport management fees by 15% or more (Table 8).

1.1. Related literature

The paper relates to different strands of the literature in finance and industrial organization. Methodologically, we develop and estimate a framework for understanding an investor's demand for brokerage services using a standard discrete choice demand model in the industrial organization literature.[9] An advantage in our setting is that we observe each investor making thousands of trades, which allows us to estimate demand for brokerage services at the individual investor level. Furthermore, due to institutional features of the market, prices are set in a quasi-exogenous manner in terms of cents per share traded. These two features make the brokerage market an ideal application for these demand estimation tools.

The paper also builds on the empirical literature on brokerage services and institutional trading patterns. Using an earlier version of our data, Goldstein et al. (2009) provide a useful description of the institutional brokerage industry, and building on their findings, we develop and estimate a framework for understanding execution decisions that allows us to quantify how broker-provided services influence execution decisions and how investors value these services.

There is a growing body of work arguing that one of the primary services brokers render is

[8]In general, we study how MiFID II and unbundling impact the use of soft dollars, investment management fees, and market transparency. A related but separate interpretation of unbundling would be whether research services and transaction services are provided by the same broker. We focus on the former because it is more closely related to proposed regulations. There is also recent research examining the preliminary impacts of MiFID II on the supply of sell-side research (Fang et al. (2020); Guo and Mota (2021); Lang et al. (2019)). Evidence from Guo and Mota (2021) suggests that the implementation of MiFID II led to a 7.45% decline in research coverage in Europe. While we study a different aspect of MiFID II, the finding that research coverage falls following the implementation of MiFID II is consistent with our empirical finding that roughly 10%+ of investment managers place no value on research (Section 6.3).

[9]This methodology has been used in other financial applications such as demand for bank deposits (Dick (2008); Diamond et al. (2020); Egan et al. (2017a); Egan et al. (2017b); Wang et al. (2020); Xiao (2020); and Xiao et al. (2020)), bonds (Egan (2019)), annuities (Egan et al. (2020); Koijen and Yogo (2016)), and mortgages (Benetton (2021); Buchak et al. (2018); Jiang (2020) and Robles-Garcia (2019)).

information provision. Our work draws inspiration from recent theoretical papers that highlight the role of financial intermediaries in creating value through information production (Babus and Kondor (2018) and Glode and Opp (2016)) and the dispersion of information in markets more generally (Duffie and Manso (2007) and Duffie et al. (2015)). Previous empirical work argues that brokers are hubs of informal research, providing their clients with details regarding informed order flow (Boyarchenko et al. (2021), McNally et al. (2015); Chung and Kang (2016); Li et al. (2017); and Di Maggio et al. (2019)), ongoing fire sales (Barbon et al. (2019)), and upcoming analyst recommendations (Irvine et al. (2006)). While previous empirical work documents that brokers are a source of informal research and act on that information, we quantitatively document how investors value this information and how it impacts their trading decisions.

Our paper also relates to the literature studying the information produced by sell-side research analysts.[10] In contrast to much of the previous literature, we examine the value of sell-side research using the revealed preferences of institutional investors, the consumers of sell-side research. In line with the previous literature, we find that analysts produce valuable information and, using our structural model, quantify the premium that investors attach to that research. We uncover significant heterogeneity in the premium that investors are willing to pay for information and highlight that the venue-routing decision is a multidimensional one, where formal/informal research and implicit/explicit costs all play a significant role. We also provide insight into how investors pay for research using soft-dollars as documented in other settings (Blume (1993); Conrad et al. (2001))

2. Framework: institutional demand for brokerage services

We develop an empirical model of broker choice. Specifically, we examine an institutional investor's decision regarding *where* to execute her trade, conditional on the investor's initial decision to trade a specific security. We model an investor's execution decision as a multinomial choice problem where the investor has a trade order she needs to execute and can route her order through any of the n available brokers denoted $l = 1, ...n$. Investors choose a broker based on the associated costs and services. For convenience and consistent with the literature on demand estimation, we initially write the investor's problem in terms of utility maximization but show below that this translates into the investor's profit maximization/cost minimization problem.

The expected indirect utility derived by investor i of executing trade j in industry sector k through brokerage firm l at time t is given by:[11]

$$E[u_{ijklt}] = -\alpha_i f_{ijklt} + X'_{klt}\beta_i + \mu_{ilt} + \xi_{iklt} + \epsilon_{ijklt}. \quad (1)$$

Investors pay an investor-trade-broker-sector specific commission/broker fee f_{ijklt} for executing a

[10] There is a broad literature documenting that research analysts produce valuable information including: Womack (1996), Barber et al. (2001), Barber et al. (2003), Jegadeesh et al. (2004), Birru et al. (2019), and Bharath and Bonini (2019).

[11] We focus on an investor's expected utility of trading with a particular broker, as opposed to realized utility, because the investor may not perfectly observe all of the relevant characteristics, such as realized price impact, before trade execution.

trade through broker l, from which she derives dis-utility $\alpha_i f_{ijklt}$. The parameter $\alpha_i > 0$ measures the investor's sensitivity to brokerage fees. Note that the parameter α_i varies across investors which implies that investors have potentially different elasticities of demand. We allow broker fees f_{ijklt} to potentially vary across investors, brokers, and trades, although previous research (Goldstein et al. (2009)) suggests that, while brokers charge different investors different fees, fees within a broker-investor pair appear to sticky. As discussed in Section 4, broker fees are typically negotiated infrequently and are used to compensate brokers for the bundle of services they provide.

Investors also derive utility from other brokerage services captured in the term $X'_{klt}\beta_i + \mu_{ilt} + \xi_{iklt} + \epsilon_{ijklt}$. The vector X_{klt} is a vector of broker-specific characteristics that reflect differences in execution services, such as price impact, speed, access to dark pools, and/or information. For example, some brokers may have more skilled traders than other firms and consequently provide better trade execution resulting in a lower transaction price (i.e., lower price impact). Furthermore, trading ability may vary within a brokerage firm across different securities and over time. For example, Goldman Sachs could provide better execution for stocks in the technology sector, while Morgan Stanley provides better execution for stocks in the financial sector. Our framework allows for such differences. The vector X_{klt} also captures the quality of research and other information provided by the brokerage firms. Arguably, investors allocate trades to brokers taking into consideration the research and other services that the investor can receive from the broker once a stable relationship is established similar to the framework proposed and studied in Goldstein et al. (2009). For example, Goldman Sachs may offer better research coverage or be privy to better information regarding stocks in the technology sector than Goldman Sachs's competitors, and an investor is likely to internalize these dimensions. The vector β_i reflects investor i's preferences over the broker characteristics X_{klt}. We again allow preferences for the various brokerage services captured in X_{klt} to vary across investors. Some investors may place a higher value on sell-side research while others place a higher value on execution.

Brokerage firms may differ in their quality of services along other dimensions beyond those captured in X_{klt}. For example, some brokerage firms may have access to their own proprietary algorithms and technology. The term μ_{ilt} is an investor-by-broker-by-time fixed effect that captures these broad differences in technology across brokerage firms. Note that this broker fixed effect (μ_{ilt}) varies across time to capture broker-specific changes in technology (i.e., the addition of new algorithm). The broker fixed effect also varies across investors to allow investors to have heterogeneous preferences over the different brokers. Because these broker fixed effects vary at the investor-by-time level they also capture specific bilateral relationships between investors and brokers.

The indirect utility includes two unobservable terms ξ_{iklt} and ϵ_{ijklt}. The term ξ_{iklt} is a time-varying investor-by-broker-by-sector latent variable that measures a brokerage firm's execution services in ways not captured by X_{klt} or μ_{ilt}. For example, Goldman Sachs's ability to efficiently trade a stock may vary over time in a way that is not captured in the vector X_{klt} or μ_{ilt}. Alternatively, an investor may prefer to trade with Goldman Sachs in the mining sector because it has a long-standing relationship with Goldman's head trader in the mining sector. This unobservable term also captures the correlation in a firm's trading patterns across trades, sectors, and over time. For example, a high value of ξ_{ijklt} helps explain why an investor i executes a larger than expected, given observable covariates, number of trades with broker l in a given market k and period t. In our empirical analysis, we do not need to make any distributional assumptions over the investor-by-broker-by-sector error terms, and it may be natural to think they are correlated across sectors and over time.

Lastly, the variable ϵ_{ijklt} reflects an investor-by-trade-by-broker-by-sector-by-time, latent, demand/profit shock. The term ϵ_{ijklt} captures preference heterogeneity within an investor across different trade ideas. For example, an investor may prefer to route a particular trade in the financial sector to Goldman Sachs while routing other trades in the financial sector to Morgan Stanley. The term ϵ_{ijklt} also potentially captures an investor's time-varying expectations about the quality of services a broker offers not captured in the vector X_{klt}. The parameter ϵ_{ijklt} introduces additional heterogeneity to help explain why we see a given investor trade through multiple brokers at the same time in a given sector. We can therefore write an investor $i's$ expected indirect utility of executing trade idea j in sector k with broker l at time t in terms of the trade-specific (ϵ_{ijkt}) and non-trade-specific, average, utility component $(\overline{u}_{iklt})$:

$$E[u_{ijklt}] = \overline{u}_{iklt} + \epsilon_{ijklt} \tag{2}$$

where $\overline{u}_{iklt} = -\alpha_i f_{iklt} + X'_{klt}\beta_i + \mu_{ilt} + \xi_{iklt}$.

The units of eq. (1) are in terms of utils; however, by scaling eq. (1) we can interpret each coefficient in the utility function in terms of expected profits:

$$E[\pi_{ijklt}] = -f_{ilkt} + \frac{1}{\alpha_i}X'_{klt}\beta_i + \frac{1}{\alpha_i}\mu_{ilt} + \frac{1}{\alpha_i}\xi_{iklt} + \frac{1}{\alpha_i}\epsilon_{ijklt}. \tag{3}$$

The vector β_i/α_i captures how the various services offered by a brokerage firm translate into an investor's profits. For example, the coefficient corresponding to research, $\beta_i^{Research}/\alpha_i$, tells us how investors value research services offered by brokerage firms in terms of the present value of expected future profits.

Investors choose the brokerage firm in the set $\mathcal{L} = \{1, 2, ...n\}$ that maximizes the investor's expected utility

$$\max_{l \in \mathcal{L}} E[u_{ijklt}]. \tag{4}$$

Under the assumption that the investor-by-trade-by-broker-by-sector-by-time specific profit shock, ϵ_{ijklt}, is distributed i.i.d. Type 1 Extreme Value,[12] as is standard in the multinomial choice literature,

[12]Note that the model includes two latent terms ξ_{iklt} and ϵ_{ijklt}. The term ξ_{iklt} captures the general unobserved correlation in a firm's trading patterns within and across markets and over time. Thus, the assumption that ϵ_{ijkt}, conditional on ξ_{iklt}, is i.i.d. across trades is relatively benign.

the probability that investor i executes her trade through firm l is given by

$$\Pr(l) = \frac{\exp\left(-\alpha_i f_{iklt} + X'_{klt}\beta_i + \mu_{ilt} + \xi_{iklt}\right)}{\sum_{m\in\mathcal{L}} \exp\left(-\alpha_i f_{ikmt} + X'_{kmt}\beta_i + \mu_{imt} + \xi_{ikmt}\right)}. \tag{5}$$

The above likelihood corresponds to the multinomial logit distribution and is the core of our estimation strategy below. Estimation of this demand framework is straightforward, and it allows us to measure how institutional investors trade-off broker-provided services. We describe estimation in Section 4.

Lastly, while we cast our framework in the context of an investor's decision regarding where to execute her trade *conditional* on the initial decision to trade a specific security, the model and corresponding estimates also generalize to the setting where brokers influence an investor's initial decision of whether or not to trade. One might think that the services offered by a brokerage firm and the expected profits of trading with a particular brokerage firm could induce an investor to make additional trades. For example, broker research could motivate an investor to trade. In our baseline framework, an investor needs to trade a security and chooses among n brokers to execute the trade through. Without any loss in generality, one could recast our model to include, in addition to choosing where to trade a security among n brokers, the outside option of not trading the particular security (which can also be influenced by brokers). As discussed below in our estimation section and further in Appendix B, adding the outside option of not trading produces numerically equivalent estimation results.

3. Data

3.1. Ancerno data

We use information about institutional transactions from Abel Noser Solutions, formerly Ancerno Ltd. (the name 'Ancerno' is commonly retained for this data set). The company performs transaction cost analysis for institutional investors and makes the data available for academic research under the agreement of non-disclosure of institutional identity.

We have access to data covering the period from 2001 to 2014. The data set consists of over 300 million trades. For each trade, we observe the names of the parties involved (broker and investment manager), the security traded, execution price, and the fee. We restrict our attention to those observations where we observe complete trade information (parties involved, security, date, and broker fee) where the investor reported paying a fee to the broker.[13] Importantly, we aggregate child orders into a unique parent order by summing all the shares traded by the same institution on the same day/stock/side of the market with the same broker. This choice overcomes the potential criticism that child orders cannot be considered as independent observations. We also focus our

[13]We drop observations where the investor does not report paying a positive broker fee to the broker. We drop these trades because we do not observe whether these zero-fee trades are indeed zero-fee trades or simply observations with missing fee data. In untabulated results, we re-estimate our baseline demand specifications where we include these trades and find comparable estimates.

attention on those institutional investors that made at least 1,000 trades in the data set. The final data set covers 383 investment managers trading across 1,510 different brokers.

Previous literature has established the merits of this data set (see Hu, Jo, Wang, and Xie (2018) for a detailed description of the structure and coverage of the data). First, clients submit this information to obtain objective evaluations of their trading costs rather than to advertise their performance, suggesting that the data should not suffer from self-reporting bias. Second, Ancerno collects trade-level information directly from hedge funds and mutual funds when these use Ancerno for transaction cost analysis. It's worth noting that pension funds may instruct the managers in whom they have invested to release their trading activities to Ancerno as part of their fiduciary obligations under ERISA regulation. Third, Ancerno is free of survivorship biases as it includes information about institutions that at some point terminated their relationship with Ancerno.

Previous studies, such as Puckett and Yan (2011), Anand, Irvine, Puckett, and Venkataraman (2012, 2013), have shown that the characteristics of stocks traded and held by Ancerno clients as well as the respective returns are comparable to those reported in mandatory 13F filings. Estimates suggest that trades recorded in Ancerno account for 10% to 19% of all institutional trading volume in the U.S. stock market (Puckett and Yan (2011), Hu, Jo, Wang, and Xie (2018)). The data is organized at different levels; at the trade level, we know the transaction date and time at the minute precision, the execution price, the number of shares that are traded, the side (buy or sell), and the stock CUSIP.

3.2. Equity research data

To help examine the different factors driving an investor's execution choice, we match our trade-level Ancerno data to sell-side equity research data from Thomson Reuters I/B/E/S and Institutional Investor. Thomson Reuters I/B/E/S is a database that provides equity analyst recommendations. We use the I/B/E/S data to determine each brokerage firm's analyst coverage for each sector over time. We merge our trade-level data with the I/B/E/S equity analyst recommendations at the brokerage firm, by year, by industry (GICS 6 Industry Code) level.[14] Table 1 displays the corresponding summary statistics. The key variable of interest is the number of analysts employed by a brokerage firm in a given sector. The average brokerage firm employs 1.48 analysts in a given sector.

We also merge our trade-level data with analyst data from Institutional Investor. Each year, Institutional Investor publishes an "All-America Research Team" where it ranks the top three equity analysts in a given sector for that year. We use the Institutional Investor data to determine the number of top-rated analysts employed by each brokerage firm in each sector and year. We merge our trade-level data with the All-American Research Team data at the year-by-sector-by-brokerage-firm level. Table 1 displays the corresponding summary statistics. The average brokerage firm in our sample employs 0.17 top analysts in a given sector and year. Previous work has shown that

[14]We merge the I/B/E/S analyst data to the brokerage firm names using data from FINRA's BrokerCheck website and a leading social networking website. As described below, FINRA's BrokerCheck data provides data, including the employment history, on the universe of individuals registered in the securities industry, including equity research analysts.

these top analysts provide more accurate forecasts (Stickel (1992)). Evidence from the brokerage industry indicates that these types of industry polls are critical for the evaluation and careers of research analysts (Groysberg and Healy (2013)). The purported policy at Lehman Brothers was for its research analysts to make "Institutional Investor or die" (Nanda et al. (2008)). These variables help capture the quality of research services at the year-by-sector-by-brokerage-firm level.

3.3. BrokerCheck data

We also examine how execution varies with the quality of a firm's traders. We merge our trade-level data with equity trader data from BrokerCheck. The Financial Industry Regulatory Authority (FINRA) maintains the website BrokerCheck which contains employment, qualification, and disclosure history for the universe of registered securities representatives over the past ten years. Our data covers the universe of registered securities representatives over the period 2005-2018 as described further in Egan et al. (2019).

As of 2017, there were 18,000 actively registered individuals licensed to trade equities in the U.S.[15] The average trader in our sample has 12 years of experience in the industry. FINRA also requires that registered representatives report any customer disputes, regulatory offenses, and/or criminal offenses. We examine whether the traders in our sample have engaged in misconduct, where misconduct is defined as per Egan et al. (2019) as any customer disputes that resulted in a settlement/award, regulatory offenses, criminal offenses, and/or terminations for cause. Roughly 6.50% of the equity traders in our sample have a record of misconduct. Table 1 indicates that at the average brokerage firm in our sample, roughly 0.20% of the traders received a misconduct-related disclosure in a given year.

Although we observe the identities of each trader, we do not observe the specific securities they trade. Consequently, we merge the BrokerCheck equities trader data with our Ancerno trade-level data at the brokerage firm-by-year level. In our analysis, we examine how much investment managers value various characteristics of a brokerage firm, including the number of traders at the firm, average trader experience, and the percent of traders previously reported for misconduct. Using BrokerCheck data, we are also able to determine the physical office locations of the brokerage firm traders and many of the investors in our data set. We calculate the physical distance in miles between each broker-investor pair, using the modal zip code of a broker's equity traders and the modal zip code of the investor's employees that are registered with FINRA. While the average distance between an investor and a broker in our sample is 668 miles, 33% of our broker-investor trading pairs are within 100 miles of each other.

[15]We determine which individuals in BrokerCheck are equity traders based on whether or not the individual has a Series 55 license. The Series 55 license, known as the Equity Trader Qualification License, entitles an individual to participate in equity trading. The

4. Estimation

We use the Ancerno micro transaction-level data to estimate the broker choice/demand model from Section 2. The model is straightforward to take to the data and allows us to determine how investors value the services that brokerage firms provide. Our estimation procedure most closely follows that of Berry (1994) and Berry et al. (1995). However, the extensive and detailed nature of the data allows for a rich flexible estimation procedure where we are able to estimate the Berry (1994) model at the investor-level. We observe tens of thousands of choices for each individual investor which allows us to flexibly recover the individual preferences of each investor without imposing any assumptions over the distribution of investor preferences α and β.

4.1. Empirical framework

Following the framework from Section 2, the share of trades investor i executes with broker l in sector k at time t is can be written as

$$s_{iklt} = \frac{\exp\left(-\alpha_i f_{iklt} + X'_{klt}\beta_i + \mu_{ilt} + \xi_{iklt}\right)}{\sum_{m\in\mathcal{L}} \exp\left(-\alpha_i f_{ikmt} + X'_{kmt}\beta_i + \mu_{imt} + \xi_{ikmt}\right)} \tag{6}$$

We define market shares at the investor-by-month-by-sector level, such that s_{ilkt} reflects the dollar value of transactions (share price $\times$ quantity) executed through broker l in sector k and in month t relative to the dollar value of all other transactions made by investor i in sector k and in month t.[16] In Appendix A.1, we also try different market definitions, using higher levels of aggregation across sectors and over longer periods (annually) to reflect the fact that execution decisions may be made at lower frequencies. Overall, we find that the choice of aggregation has little impact on the parameter estimates (Table A1).

Following Berry (1994), we can rewrite the market share of broker l in a given market (month-by-investor-by-sector) as

$$\ln s_{iklt} = -\alpha_i f_{iklt} + X'_{klt}\beta_i + \mu_{ilt} + \xi_{iklt} - \ln\left(\sum_{m\in\mathcal{L}} \exp\left(-\alpha_i f_{ikmt} + X'_{kmt}\beta_i + \mu_{imt} + \xi_{ikmt}\right)\right) \tag{7}$$

Notice that the non-linear term $\ln\left(\sum_{m\in\mathcal{L}} \exp\left(-\alpha_i f_{ikmt} + X'_{kmt}\beta_i + \mu_{imt} + \xi_{ikmt}\right)\right)$ is constant in a given market. Therefore we can estimate eq. (7) using linear regression where we include an investor-by-sector-by-time market fixed effect (μ_{ikt}) to absorb the non-linear term.[17] We estimate

[16] We define the market at the investor-by-month-by-sector level rather than at the investor-by-month-by-stock level to match how brokerages are organized. For example, sell-side research teams are typically organized at the sector level. Aggregation helps facilitate estimation and allows us to estimate the model using linear regression rather than maximum likelihood or other non-linear estimation methods.

[17] Note that the nonlinear term captures the utility that the investor derives from trading with any other potential trading partner in his/her choice set. Because this term is absorbed by the fixed effects, we do not need to observe or even define an investor's full choice set. Consequently, as discussed further in Appendix B, if we were to re-estimate our model from Section 2 where investors have the option of not trading, the estimates would be numerically equivalent to our baseline estimates.

the linear specification

$$\ln s_{iklt} = -\alpha_i f_{iklt} + X'_{klt}\beta_i + \mu_{ilt} + \mu_{ikt} + \xi_{iklt} \tag{8}$$

where X_{klt} is our vector of broker-by-sector-by-time characteristics and μ_{ilt} is an investor-by-broker-by-time fixed effect. We describe the construction and details of each of our broker characteristics X_{klt} in the proceeding section. In our baseline specifications, (Section 5), we restrict the preferences of investors to be the same across investors such that $\alpha_i = \alpha$, $\beta_i = \beta$, and $\mu_{ilt} = \mu_{lt}$. However, we relax this assumption in Section 6.3 where we allow preferences to vary flexibly across investors.

We micro-found the demand system in Section 2. Micro-founding the demand system provides additional interpretation and allows us to investigate counterfactuals in Section 7. However, it is worth noting that our estimates also have a reduced-form interpretation in addition to a structural interpretation; we are essentially regressing broker trade volumes on a vector of broker characteristics. Thus, our estimation results are more general than what the model in Section 2 entails.

4.2. Broker characteristics

We are interested in the factors that drive institutional investors' execution decisions across brokers. Using our rich data set described in Section 3 we analyze how fees, research, quality of execution, and information drive investor decisions. Here, we provide a description of each variable, its measurement, and how we incorporate the variable in our estimation strategy. We measure each variable on a trade-by-trade basis, and then aggregate each variable at the broker-investor-sector-month level for estimation.

4.2.1. Explicit trading costs (broker fees):

Brokers typically charge investors a fee for each share of stock traded. We measure the broker fees paid on a per-trade basis as the total fee paid relative to the value of the transaction.

$$f_{ijklt} = \frac{Total\,Fee\,in\,USD_{ijklt}}{Value\,of\,Transaction\,in\,USD_{ijklt}} \tag{9}$$

The average transaction fee is 13 basis points (bp). Figure 1a displays the distribution of broker fees paid by investors. There is substantial variation in fees paid by investors. The standard deviation of fees is 13bps and fees range from near zero to upwards of 20bps. The average mutual fund turned over 54% of its portfolio in a given year over the period 2001-2014, which suggests that the variation in trading fees could be costly on an annual basis. For the average mutual fund, a one-standard-deviation increase in broker fees translates to an annualized cost of 14bps ($\approx 2 \times 54\% \times 13$bp) relative to the fund's total assets.[18] To put these numbers in perspective, the average mutual fund over that same period charged an expense ratio of 0.87% (2018 Investment Company Factbook).

[18]When calculating annual trading costs, we multiply turnover by two to account for the fact that turning over a portfolio involves both a buy and a sell trade.

Evidence from Goldstein et al. (2009) suggests that broker fees are negotiated at the broker-investor level and compensate brokers for the full bundle of services brokers provide in addition to execution. Goldstein et al. (2009) finds that at the broker-investor level fees appear sticky in terms of cents per share, which suggests that fees are negotiated infrequently. Regardless, one of the standard issues in demand estimation that we need to address is the potential endogeneity of broker fees. Fees are potentially endogenous if brokers observe demand shocks, ξ_{iklt},before setting their fees. Conceptually, the idea is the following: if brokers know that their services are in high demand and/or anticipate high order flow, they may adjust their fees accordingly. In general, this potential endogeneity problem will bias the OLS estimates of $-\alpha$ upwards such that we would underestimate an investor's responsiveness to fees.

While evidence from Goldstein et al. (2009) suggests that, by and large, brokers do not update their fees on a transaction-by-transaction basis, we address the endogeneity problem using instrumental variables. A unique feature of the institutional setting is that most brokerage firms charge investors a fixed dollar amount per shares of stock traded (see Goldstein et al. (2009)). Figure 1b displays the distribution of broker fees charged on a per-share basis. As illustrated in the figure, the fees are bunched around the whole numbers in terms of cents per share ranging between 1 cent and 6 cents per share (the mode is 5 cents per share). However, the relevant metric for a profit-maximizing investor is measuring fees in percentage terms relative to the value of a transaction. For instance, a one-cent increase in the fee per share is more costly when an investor is trading a stock priced at \$1 per share than when she is trading a stock priced at \$1,000 per share. We exploit variation in the underlying share price as an instrument for fees.

We exploit the institutional fee setting feature of the brokerage industry to construct an instrument for broker fees. We construct our instrument at the trade level as the inverse of the corresponding equity share price scaled by the average cents per share fee charged by brokerage firm l:

$$IV_{ijklt} = \frac{1}{Share\,Price_{jt}} \times \overline{Fee\,Per\,Share\,In\,USD_l}. \tag{10}$$

The key variation in the data is that relative fee differences between brokers, in terms of cents per share, are magnified when the underlying share price is low. Thus the instrument is correlated with our measure of fees in percentage terms f_{ijklt} because, all else equal, a decrease in the share price makes the fixed per-share fee more expensive on a relative basis. What matters for the relevancy condition is that broker fees (i) are set on a per-share basis, as it is common in the industry, and that (ii) brokers do not immediately update their fees, in terms of cents per share, in response to share price movements. Note that our instrumental variables strategy does not rely on the fact that broker's fees are typically rounded to the nearest cent, although the rounding does provide additional potentially exogenous variation in broker fees. As discussed in the next section, our instrument yields Cragg-Donald F Statistics well above 100 in each specification (Cragg and Donald (1993)), suggesting a strong first-stage significance of the instrument. The instrument satisfies the exogeneity condition essentially as long as stock price movements are orthogonal to the investor-broker-market-time specific demand shocks ξ_{iklt}. While movement in stock prices would certainly be correlated with an investor's decision to trade, what matters for our setting is that movements in stock prices are not correlated with *who* an investor trades with at a particular moment in time. Recall that our regression specifications include broker-time and investor-sector-time fixed effects; thus the exogeneity condition requires that the share prices are uncorrelated with time-varying quality differences across brokers.

4.2.2. Formal research:

We measure the level and quality of a brokerage firm's research coverage in a particular sector along two dimensions using our I/B/E/S and Institutional Investor data sets. First, we include the number of analysts a brokerage firm employs in a given sector and year.[19] Second, we control for the number of top analysts as reported by Institutional Investor that the brokerage firm employs in a given sector and year. We examine whether investors are more likely to trade through brokers who have analyst coverage in the corresponding sector and measure the value that investors place on those sell-side analysts.

As discussed previously, we are worried about the potential endogeneity of broker characteristics such as fees. It is possible that, just as with fees, firms adjust their research coverage in response to demand shocks (ξ_{iklt}), which would cause research coverage to be endogenous. For example, if a broker anticipates a high demand shock, the broker might find it optimal to reduce its research coverage, which would typically cause the corresponding OLS estimated coefficients to be biased downwards. The reason why researchers in the demand estimation literature to date have specifically been more concerned about the endogeneity of prices (broker fees f in our setting) rather than other product characteristics (broker characteristics X in our setting) is that prices (fees) are likely to be the margin of adjustment in response to time- and sector-varying demand shocks. For example, in practice, it might be difficult for a broker to endogenously change its research coverage in the short-run (e.g., on a month-by-month basis) because hiring research analysts is a lengthy and involved process that regularly takes a year (Groysberg and Healy (2013)). Consistent with this intuition, we find that research coverage is quite sticky in the data. The 12-month autocorrelation of the number of analysts a brokerage firm employs in a given sector is 0.88, and the autocorrelation in of the number of top analysts is 0.79 (Table A3). Consequently, we believe it less likely, although possible, that firms are updating their research coverage in response to demand shocks. As a robustness check, in the Appendix A.3 we run additional specifications where we instrument for research coverage using historical coverage and find quantitatively similar results (Table A4).[20]

[19]To account for outliers we winsorize the number of analysts at the 1% level.

[20]The rationale behind using lagged research coverage as an instrument for current research coverage is that it is likely to be relevant because it's costly to adjust research coverage and the costs of providing research analyst coverage are correlated over time, and it is potentially exogenous because lagged research coverage is likely to be uncorrelated with current demand shocks. The endogeneity concern would be that if demand shocks are highly serially correlated and research coverage is endogenous, then lagged research coverage could be correlated with current demand shocks.

4.2.3. *Informal research:*

Recent evidence has highlighted the role played by financial intermediaries in creating value through information production (Babus and Kondor, 2018; Barbon et al. 2019). Brokers may have access to additional information about market conditions, trends, and specific stocks. We collectively label this information as 'informal research'. We use two different measures to capture a broker's ability to produce informal research based on Di Maggio et al. (2019). First, we calculate the eigenvector centrality of the broker in the network where we define the network at the sector-by-month level. The eigenvector centrality measure takes into account all direct trading partners (i.e., investors) as well as indirect trading partners (i.e., other brokers and investors) and is computed by assigning scores to all brokers in the network.[21] What counts is not only the number of connections of a broker but *who* the broker is connected to. All else equal, being connected to a more central manager leads to a higher centrality score for the broker. We construct eigenvector centrality at the sector-by-month level for each investor i and broker l pair, $Eigenvector\_Centrality_{iklt}$. To avoid clear endogeneity concerns, we remove all of investor $i's$ trades from the network when computing the centrality of broker l in sector k at time t, $Eigenvector\_Centrality_{iklt}$.

We also control for whether a broker is "informed" in a given market following Di Maggio et al. (2019). The authors find evidence suggesting that after executing an "informed" trade, brokers tend to share that information with other investors. Following these authors, we define an "informed trade" as an abnormally large (75th percentile) profitable trade made by a hedge fund. Roughly 1.7% of the trades in our sample are classified as informed. We measure whether a broker received an informed trade at the month-by-sector level , $Informed_{klt}$. In the analysis, we control for $Informed_{klt-1}$ (i.e.,$Informed_{klt}$ lagged by one month) to measure how informed order flow influences the proceeding execution decisions of other investors.[22]

4.2.4. *Implicit trading costs (price impact):*

Implicit trading costs may arguably be just as important as explicit trading costs. Anand et al. (2012) show that brokers differ in their execution quality persistently. We measure the implicit trading cost of a trade using the implementation shortfall (Perold (1988), Wagner and Edwards (1993)). As described in Anand et al. (2012), the execution shortfall reflects the bid-ask spread, the market impact, and the drift in price. With this in mind, we call this variable price impact and define it as the stock price at which the trade was ultimately executed relative to the stock price at

[21]Note that we are modeling a bipartite network. One set of nodes contains brokers, the other set contains investors. Direct connections within each set are not possible. Connections can only occur across the two sets. That is, brokers are not directly connected to each other in the trading network and are only connected to each other indirectly through other investors in the trading network.

[22]By construction, the variable $Informed_{klt}$ indicates that one manager executed an informed trade through broker l in sector k in month t. Thus $Informed_{klt}$ will be, at least partially, mechanically related to the trades executed through a broker. Consequently, we lag $Informed$ by one month, to proxy for how the execution of informed order flow influences the proceeding execution decisions of other investors.

the time the order was placed,

$$Price\,Impact_{ijklt} = \left(\frac{Execution\,Price_{ijklt} - Benchmark\,Price_{ijklt}}{Placement\,Price_{ijklt}}\right) \times Side_{ijklt}. \tag{11}$$

The variable $Side_{ijkt}$ is equal to 1 if the trade is a buy trade and equal to -1 if the trade is a sell trade. All else equal, investors prefer a lower price impact, and a high price impact is indicative of worse execution.

To calculate the price impact in our data, we first calculate the weighted-average price impact at the broker-by-month-by-stock level to construct the variable $Price\,Impact_{lst}$, where l indexes the broker, s indexes the stock, and t indexes the month. To account for time-varying differences in the liquidity of different stocks, we residualize the variable $Price\,Impact_{lst}$ on a vector of stock-by-month fixed effects to construct the variable $Price\,Impact^*_{lst}$. This is similar to the way Anand et al. (2012) measure trading desk performance, where they regress price impact on a vector of stock-specific characteristics. Lastly, we calculate the weighted average of $Price\,Impact^*_{lst}$. at the broker-by-sector-by-month level ($Price\,Impact^*_{lkt}$), which corresponds to our definition of a market and is the primary observational unit of our analysis. The variable $Price\,Impact^*_{lkt}$ measures a broker's trading ability at the sector-by-month level.

There are a handful of potential concerns with our price impact measure $Price\,Impact^*_{lkt}$ that merit further discussion. First, it is inevitably measured with noise. Our empirical measure of price impact reflects both the true variation coming from price impact as well as variation from changes in the underlying fundamentals from the stock. For example, even if markets were perfectly liquid, we would expect the execution price to potentially differ from the placement price due to the variation in fundamentals. This type of measurement error will potentially cause our estimates to suffer from attenuation bias.

Second, we are using contemporaneous price impact as a control variable which includes information unavailable to investors at time t. Ideally, we would like to be able to control for an investor's expectations about the price impact at time t, given the investor's information set at time $t-1$, $E[Price\,Impact^*_{lkt}|\mathcal{I}_{t-1}]$.

Lastly, and related to the previous point, $Price\,Impact^*_{lkt}$ could suffer from reverse causality. If a broker experiences a positive demand shock in a specific sector such that a large number of investors choose to trade through the broker, this could lead to the broker providing either better or worse execution due to increased trading volumes. To address these issues we use both contemporaneous and lagged price impact as proxies for an investor's price impact expectations:

$$Price\,Impact^*_{lkt} = E[Price\,Impact^*_{lkt}|\mathcal{I}_{t-1}] + \eta_{ijklt} \tag{12}$$

$$\overline{Price\,Impact^*_{lkt-12}} = E[Price\,Impact^*_{lkt}|\mathcal{I}_{t-1}] + \nu_{ijklt} \tag{13}$$

where $\overline{Price\,Impact^*_{lkt-12}}$ is the lagged twelve-month rolling weighted average of broker $l's$ price impact in sector k. We then use contemporaneous price impact as a proxy for investor price impact expectations and construct lagged price impact as an instrument. Previous work finds that there is strong persistence in broker trading performance (Anand et al. (2012)) which indicates that our instrument will be relevant (i.e., there are systematic differences across brokers that determine their execution quality). To the extent that the measurement error η_{ijklt} is orthogonal to ν_{ijklt}, then using instrumental variables will help address the potential measurement error issues with our proxies for price impact.

As an additional robustness check, we examine investors' sensitivity to large trading costs. Specifically, in Appendix A.4 we re-estimate our baseline specification where we control for implicit trading costs with the variable $LargePriceImpact_{lkt}$ which measures whether the price impact was greater than 0.25% (roughly the 50th percentile). One might expect that investors avoid brokers with a track record of particularly poor execution, which may be more salient and predictable than average execution.

4.2.5. Broker fixed effects:

In our main specifications, we include broker (μ_l) fixed effects, broker-by-time fixed effects (μ_{lt}), and, in our most stringent specification, broker-by-time-by-investor fixed effects (μ_{ilt}). The broker-by-time fixed effects capture broad, potentially time-varying, differences across brokerage firms. For example, some brokerage firms may have better algorithms and technology. These differences in trading technologies across firms will be captured in our broker fixed effect.

The broker-by-time-by-investor fixed effects help capture potential long-term and time-varying relationships between brokers and investors. For example, an investor may prefer to trade with a specific set of brokers based on the investor's past relationships and negotiations with these brokerage firms. Also, because we allow broker fixed effects to vary across investors, these fixed effects capture each investor's subjective preferences over the brokers in our sample.

5. Results

Table 2 presents our main sets of estimation results corresponding to eq. (8). The columns differ for the set of fixed effects and whether or not we estimate the model using ordinary least squares or instrumental variables. In column (1) we report our baseline set of results where we estimate the model using ordinary least squares and include market fixed effects. In column (2) we re-estimate our baseline model instrumenting both fees and expected price impact as described in Section 4. Lastly, in columns (3)-(5) we include broker, broker$\times$time, and broker$\times$time$\times$investor fixed effects to capture differences in trading service quality across brokerage firms and control for the specific relationships between investors and brokerage firms. In the proceeding subsections, we discuss and interpret how investors respond and value each of the brokerage firm characteristics.

5.1. *Baseline Results*

5.1.1. *Fee sensitivity:*

One of the primary coefficients of interest is how sensitive institutional investors are to fees. In each column, we estimate a negative and significant relationship between trading volumes and brokerage fees. As expected, the estimated effect becomes more negative once we employ instrumental variables. We would expect the OLS estimated fee coefficient to be biased upwards due to the endogeneity of fees. If brokers anticipate a positive demand shock (ξ_{iklt}), they will find it optimal to charge a higher fee. Thus, $-\alpha$ will be biased downwards. The first stage of our instrumental variables is quite strong. We report the corresponding Cragg-Donald F Statistic at the bottom of Table 2 (Cragg and Donald (1993)). The corresponding F-statistics are above 1,000 which is substantially greater than the typical rule-of-thumb threshold (10) and the critical values for a weak instrument reported in Stock and Yogo (2005).[23]

In the bottom panel of Table 2, we interpret the estimated coefficients in terms of elasticities. In our demand framework, the investor's elasticity of demand in a given market is $\alpha(1 - s_{iklt})f_{iklt}$.[24] Consistently across our main specifications, we find evidence suggesting that demand for brokerage services is relatively inelastic, with an elasticity of roughly 0.47. The estimates imply that if a broker increases the fee it charges by 1%, its market share will decrease by an associated 0.47%.[25] One caveat is that the level of fees is low in the sample, which implies a 1% change in fees is quite small in level terms. Consequently, interpreting the coefficient α in terms of semi-elasticities or a level change in fees is perhaps more useful than interpreting α in terms of an elasticity, which captures percentage changes. Our estimates imply that if a broker were to decrease the fee it charges to an investor in a market by one standard deviation (13bp), the share of the investor's trades that the investor executes through that broker in that market will increase by 5 percentage points on average. Given that a broker's average market share in any given market is 10%, this implies that a one standard deviation increase in fees would translate to a 50% increase in trading volume.

While investors respond to fees, we still find that demand is inelastic. This indicates that there is a fair amount of unobservable product differentiation in this market. This is also consistent with the idea that broker-investor relationships are important, sticky, and costly to form/build, which creates switching costs across brokers. As shown in Table 2 and described further below, broker fixed effects and investor-broker-time fixed effects explain substantial variation in the data. Another related reason why the demand elasticity is low is that investors may prefer trading with a handful of brokers, rather than the lowest cost broker, to help maintain strong relationships and to conceal their trades. To the extent that investors wish to execute their trades through multiple trusted

[23]Stock and Yogo (2005) provide the critical values a weak instrument test for the maximal size (10%) of a 5% Wald test of $\beta = \beta_0$. The corresponding critical value with two endogenous regressors and two instruments is 7.03.

[24]The elasticity of demand is given by $\frac{\partial s_{iklt}}{\partial f_{iklt}} \times \frac{f_{iklt}}{s_{iklt}}$. Given the empirical framework, it is straightforward to show that $\frac{\partial s_{iklt}}{\partial f_{iklt}} = \alpha s_{iklt}(1 - s_{iklt})$ following eq. (6).

[25]As a point of reference, the estimated demand elasticity is similar to what other researchers have found among banking depositors. Egan et al. (2017a) estimate that the elasticity of demand for deposits is roughly 0.2-0.6 and Xiao (2020) estimates that the elasticity of demand is 0.75.

counterparties, with whom they have an established relationship, to try to better their positions, the idiosyncratic component of demand (ϵ_{ijklt}) will be important. Consistent with this idea we find in Section 6.3 that hedge fund investors, who likely have the greatest incentive to hide their trades, tend to have the most inelastic demand.

5.1.2. Value of formal research:

Most "high-touch" brokers try to attract clients' order flow by providing research services separate from trade execution. The average brokerage firm in our sample employs roughly 1.5 research analysts and 0.17 top research analysts in a given sector. The coefficient estimates in the top half of Table 2 indicate that investors are more likely to execute trades through brokers with more extensive research coverage. The coefficients on the number of research analysts and the number of top research analysts are positive and significant in each specification.

To aid in the interpretation of our results, we report the coefficient estimates in terms of willingness to pay in the bottom panel of Table 2.[26] Specifically, we scale each coefficient estimate by $\frac{1}{\alpha}$ which puts the coefficients in terms of fees rather than utils. We find that investors value sell-side research and place a premium on the top analysts ranked in Institutional Investor, which is consistent with the finding that these top-rated analysts provide more accurate forecasts (Stickel (1992)). The results in column (2) indicate that investors behave as if they are willing to pay an additional 5.33bps per trade to have access to a top equity research analyst while having access to an additional non-top analyst is worth 1.69bps. To put these numbers in perspective, the mean and standard deviation of brokerage fees is 13bps. Thus, the results in column (2) indicate that investors behave as if they are willing to pay a 40% (=5.33/13) higher fee, relative to the mean fee, to access a top equity research analyst.

One potential concern is that the number of analysts and top analysts could be proxying for some other brokerage firm characteristic. While this is indeed possible, we believe it is unlikely that are our results are completely driven by unobservable characteristics for two reasons. First, we include broker-by-month fixed effects in our most stringent specifications, so it would have to be the case that research analyst coverage is proxying for some other brokerage firm characteristic at the broker-by-sector level over time. Second, in the next section (Section 6.3) we show that investors have heterogeneous preferences over research. Our estimates indicate that those investors that we would expect to place no value on sell-side research, such as index fund managers and hedge funds, indeed place no value on sell-side research. Thus, if our results are driven by some unobserved broker-by-sector-by-investor characteristic, it would have to be that index fund managers and hedge funds also place little value on that characteristic.

[26]Note that the coefficient estimates in the top half of Table 2 correspond to a discrete choice model. Note that the marginal effect of X is given by $\frac{\partial s}{\partial x} = s(1-s)\beta$.

5.1.3. *Value of informal research:*

Our results suggest that investors are more likely to trade through brokers with access to informed order flow and brokers that that are more centrally located within the network. In each specification, the coefficients corresponding to informed order flow and eigenvector centrality are positive and significant (Table 2). To aid in interpretation, we interpret the coefficients in terms of willingness to pay (scaled by $\frac{1}{\alpha}$) in the bottom panel. The results in column (2) indicate that investors behave as if they are willing to pay an additional 2.80bps per trade to trade through a broker who has a one-standard-deviation higher centrality measure. The results are even more economically significant when we consider the informed broker measure. We find that the investors behave as if they are willing to pay an additional 2-6bps to trade through an informed broker, which is similar to and slightly higher than the value that investors place on sell-side research. Intuitively, the information that brokers provide about current order flow is potentially as important/valuable, if not more important, than the sell-side research analyst reports that are publicly released.

5.1.4. *Implicit trading costs (price impact):*

Investors pay implicit trading costs in addition to explicit fees. Our results suggest that investors are more likely to trade through brokers who provide better execution (Table 2). In columns (2)-(4) we instrument for expected price impact using lagged price impact, as described in Section 4 to account for measurement error and potential endogeneity issues. In each specification, we estimate a negative and statistically significant relationship between a broker's trading price impact and the broker's market share. We interpret the magnitudes in the bottom panel of Table 2. The results in column (2) indicate that investors behave as if they are willing to pay an additional 6bps to trade through a broker whose expected price impact is one-standard-deviation (0.64%) lower. In terms of the variation in price impact, our estimates indicate that a one standard deviation increase in price impact corresponds roughly to half a standard deviation increase in brokerage fees (0.13%). Thus, in terms of the variation of the data, the expected price impact has a first-order impact on order flows.

To the extent that expected price impact directly translates into higher execution costs, one might expect investors to trade off price impact and broker fees one-for-one. Recall that our estimate of investor's preference for price impact likely suffers from attenuation bias because our measure of price impact reflects both variation in fundamentals and the true price impact. Also, to the extent that the price impact is not perfectly observed and forecastable by investors, a Bayesian investor would find it optimal to place less weight on implicit trading costs relative to explicit costs. When we focus on large trading costs as a robustness check (Appendix A.4), we find that investors trade off implicit and explicit trading costs roughly one-for-one. This suggests that investors avoid brokers with a track record of particularly poor execution, which may be more salient and predictable than average execution.

5.1.5. *Broker fixed effects:*

We include broker fixed effects in our main specifications to measure persistent differences in unobserved quality across brokers. We find that these fixed effects explain a fair amount of variation in trading patterns. Figure 2 displays the distribution of broker fixed effects corresponding to the estimates in column (3) of Table 2. To aid interpretation we scale the broker fixed effects by the term $1/\alpha$ which converts the fixed effects from utils into the same units as fees (basis points), and the fixed effects can be interpreted in terms of an investor's willingness to pay. Note that, because all of the specifications include investor-market fixed effects, the mean fixed effect is zero by construction, such that the fixed effects measure relative differences across brokers. The standard deviation of the scaled broker fixed effects is 9 basis points (bps), which implies that investors behave as if they are willing to pay an additional 9bps in fees to trade with a broker that has a one-standard-deviation higher broker fixed effect. To put this in perspective, the standard deviation of broker fees is 13bps, which suggests that these broker fixed effects are quantitatively important for driving trading decisions.[27]

5.2. *Value decomposition*

Using the demand estimates, we can decompose the value that brokers create in terms of formal research (analyst coverage), informal research (informed order flow and eigenvector centrality), and execution quality (price impact). We calculate the value created by the broker as the utility flow associated with these observable services ($X'\hat{\beta}$), measured in utils. The value then depends on the services the broker offers X and how investors value those services β. We calculate the value of formal research as the utility flow generated by analysts and top analysts; the value of informal research as the utility flow generated by trading with informed brokers and more central brokers; and the value of execution based on our measure of $PriceImpact$.[28]

Figure 3a shows the breakdown of the value of services that brokers provide to the average institutional investor in our sample. We find that formal and informal research account for roughly

[27]In the Appendix Figure A2 we report the distribution broker-investor-time fixed effects that correspond to the estimates in column (5) of Table 2. The standard deviation of scaled investor-broker-time-fixed effects is 30bps, which indicates that broker-investor relationships are critical in this market. One caveat with interpreting the distribution of these broker-investor-time fixed effects is that many of them are imprecisely measured in the data. On average, we only observe 10 observations within a broker-investor-time triplet, which means we have limited power to precisely estimate these fixed effects. While the measurement error in our fixed effects does not impact the interpretation of our estimates in Table 2, it does impact the interpretation of the distribution of broker-investor-time fixed effects. To the extent that the measurement error is orthogonal to the true underlying distribution of fixed effects, this will cause the estimates of the fixed effects to overstate the true underlying variance of investor-broker-time fixed effects.

[28]Note that the research and information covariates ($X^{Research}, X^{Information}$) are all bounded from below by 0. Thus, any positive value research or information covariate ($X^{Research} > 0, X^{Information} > 0$) creates value for investors. Instead, the measure of execution quality, $Price\,Impact$, takes on positive and negative values. Furthermore, even though a broker may have generated a positive price impact on a trade, the broker could still have created value for the investor. This is because what matters is not the level of price impact, but the price impact relative to a counterfactual in which the investor has traded on her own or through the worst possible brokers in terms of execution quality. To calculate the value of price impact we assume that any trade with a price impact greater than 0.70% (90th percentile of the distribution) generates negative value for the investor. In other words, we measure the value of execution as $(0.70\% - Price\,Impact)\beta^{Price\,Impact}$.

50% of total value, with the former being somewhat more valuable than the latter (29.77% vs. 20.13%). The other half of value production originates from trade execution. Thus, the main conclusion from our analysis is that brokers create as much value for their clients through research and information production as they do through trade execution.

The decomposition in Figure 3a focuses on the value generated by formal research, informal research, and execution quality and ignores the value captured by the estimated broker fixed effects. Figure 3b displays an alternative version of 3a where we account for the value captured by broker fixed effects.[29] The results indicate that our broker fixed effects account for roughly 39% of the value generated by brokers, execution accounts for 31%, formal research accounts for 18%, and informal research accounts for the remaining 12%.

6. Extensions

6.1. Trader characteristics

A unique feature of our data set is that we also observe characteristics of the individual equity traders working for the brokerage firms in our Ancerno data. We can match the investor trading data from Ancerno with the trader-level data for about half of our sample.[30] We re-estimate the baseline demand specification where we control for the characteristics of each broker's traders. Specifically, we control for the number of traders a firm employs, the average experience of those traders, and whether or not those traders engage in financial misconduct.

Table 3 presents the corresponding estimates. In each specification, we estimate a negative and statistically significant relationship between trader misconduct and a broker's market share. The results in column (1) indicate that investors are indifferent between a 1pp increase in misconduct and a 0.45bps decrease in fees. Financial misconduct includes customer disputes, regulatory, and criminal offenses. These results suggest that financial misconduct costs brokerage firms money in the form of lower trading volumes.

We also find that investors prefer to trade through firms that employ more experienced traders. The results in column (2) indicate that, on average, investors are willing to pay an additional 0.70bps to trade through a firm whose traders have an additional year of experience. However, we find evidence of a non-linear relationship. Investors prefer to trade through more experienced traders up until the trader has accumulated 14 years of experience. Beyond 14 years, investors prefer to trade with less-experienced traders. This suggests that traders may learn on the job

[29] One caveat is that the estimated broker fixed effects measure relative value across brokers but not the level value. By construction, the mean broker fixed effect is zero. To pin down the level of value generated by our estimated broker fixed effects, we normalize the level of value based on the fixed effects of full-service brokerage firms, which have sell-side analyst coverage, relative to the fixed effects of non-full service/discount brokerage firms, which do not employ any sell-side analysts. We make the normalizing assumption that the average discount broker does not produce any unobservable services/value for investors. Thus, we assume that the average fixed effect of discount brokers, denoted μ^0, corresponds to zero. The average fixed effect among discount brokers is 4-5bps below the mean broker fixed effect. Given this normalization, we can calculate the level of utility/value generated by broker fixed effect μ_l as $\mu_l - \mu^0$.

[30] We can match only half of the Ancerno data set with the BrokerCheck trader-level data because BrokerCheck covers the period 2005-2018 whereas Ancerno covers the period 2001-2014.

over the first decade of their career, but their skills diminish over time. While investors appear to value the experience of the traders, we find little evidence suggesting that investors have strong preferences over the size of trading desks.

Using the trader-level data set, we can also determine the distance between investors and brokerage firms' traders for roughly 30% of the trades in our sample. We re-estimate our demand specification controlling for distance and present the corresponding estimates in Table 4. The results indicate that investors prefer to trade through brokers who are located in the same city as the investor (within 100 miles). The economic magnitude of the estimated effect is substantial. The estimates in column (2) indicate that investors behave as if they are willing to pay 10bps more per trade to trade through a broker who is located in the same city as the investor. The effect of being in the same city translates to a roughly one standard deviation decrease in brokerage fees. The effect is also somewhat surprising given that equity trades occur over the phone or electronically and not in person. These results also suggest that investors strongly prefer to trade through parties that they potentially know on a more intimate level and that relationships remain important in the industry. This is consistent with the idea that "trading is—and always has been—a relationship business."[31] Finally, we note that proximity to brokers is not capturing investor or broker location in big cities (e.g. NYC) because our specifications include broker and investor fixed effects.

6.2. Evolution over time

Our data covers the period 2001-2014, which includes important developments in the financial sector, in terms of technology, market conditions, and regulations. For example, Regulation National Market System (Reg NMS), introduced in 2008, increased the competition among exchanges, giving investment managers the opportunity to source liquidity from the exchange offering the most favorable terms. At the same time, the disruption of financial markets caused by the Great Financial Crisis may have induced managers to strengthen their ties with brokers with whom they had a reliable pre-existing relationship, as found by Di Maggio et al. (2017) for OTC markets.

To understand how these changes impacted investors' execution decisions, we re-estimate our baseline broker-choice model where we allow investor preferences over fees (α_t) and other broker characteristics (β_t) to vary from year to year. Specifically, we estimate eq. (8) where we interact broker characteristics (X) and fees (f) with year dummy variables.

We report the corresponding estimates in Table 5. Overall, the preferences of investors appear very stable year-to-year. The one exception is that investors become substantially more sensitive to explicit trading costs over time. Figure 4 displays how the average demand elasticity changed during the sample period. The elasticity of demand doubles from 0.35 to 0.70 as investment managers become more sensitive to trading costs. The fact that elasticity goes up is consistent with increasing investor attention to commissions and possibly more aggressive competition in the broker sector resulting from the rise of electronic trading and the introduction of Reg NMS. This

[31]The quote is from Johnson, Vice President of Market Structure and Technology at Greenwich Associates. [https://www.bloomberg.com/professional/blog/human-high-touch-trading-stay/] accessed 5/9/2019.

trend in elasticity is partly reversed during the Great Financial Crisis. A potential explanation for this finding is that, during that period of market turmoil, investors placed greater weight on broker-provided services and differentiation across brokers, while they cared less about fees. Additionally, during a crisis, there may be more heterogeneity in brokers' ability to provide their services. Finally, as discussed above, at a time of forced liquidations, a long-standing relationship with a broker who can reliably find counterparties for block trades is worth more than saving money on commissions.

6.3. Investor heterogeneity

In our baseline empirical analysis, we implicitly assumed that investors have the same preferences. However, in practice, different investors are likely to value broker services differently. For example, an index fund manager may be extremely price-sensitive relative to a hedge fund or active mutual fund manager. Similarly, an index fund manager would likely place no value on sell-side research while other investors may place a premium on high-quality research. An advantage of our rich empirical stetting is that we can estimate demand for these services at the investor level.

6.3.1. Estimation

We re-estimate the baseline specification (eq. 8) allowing preferences over fees (α_i) and other broker characteristics (β_i) to vary across investors. Recall from our earlier framework, that an investor's indirect utility function from trading is:

$$u_{ijklt} = -\alpha_i f_{iklt} + X'_{klt}\beta_i + \xi_{iklt} + \epsilon_{ijklt}. \tag{14}$$

In our baseline specification, we assume that preferences are constant across investors such that $\begin{pmatrix} \alpha_i \\ \beta_i \end{pmatrix} = \begin{pmatrix} \alpha \\ \beta \end{pmatrix}$. To implement our specification with heterogeneous preferences we estimate the following regression at the investor-level:

$$\ln s_{iklt} = -\alpha_i c_{iklt} + X'_{klt}\beta_i + \mu_{il} + \mu_{ikt} + \xi_{iklt}. \tag{15}$$

This allows us to recover the distribution of coefficients $\begin{pmatrix} \alpha_i \\ \beta_i \end{pmatrix}$ without placing any parametric restrictions on the distribution of coefficients. Again, observations are at the investor-by-sector-by-month-by-broker level.

To recover the distribution of investor coefficients, we separately estimate eq. (15) at the investor level such that we can recover each investor's preferences α_i and β_i. In other words, we can estimate the random-coefficients demand model using simple linear regression at the investor level. This is in sharp contrast to the way one typically has to estimate a Berry et al. (1995) (BLP) type demand system. In the standard Berry et al. (1995) set-up, the econometrician only observes aggregate demand data, rather than individual demand data. Consequently, with aggregate data, one typically has to make parametric assumptions over the distribution of preferences (α_i, β_i), and estimates the model via GMM. Estimating the model via GMM with aggregate data involves solving a non-trivial contraction mapping for each set of parameters that the econometrician searches over

(Berry et al. (1995); Nevo (2000)). Because of our unique, detailed, microdata, where we observe each investor making thousands of decisions, we can estimate our demand model using simple regression at the investor level. Furthermore, we do not need to make any parametric assumptions over the distribution of investor preferences (α_i, β_i). In most data sets, the researcher does not have enough observations at the individual level to estimate individual-specific demand functions. For power considerations, we estimate eq. (15) at the investor level where we restrict our sample to those 243 out of 383 investors that have at least 1,000 observations (sector-by-month-by-broker level).

6.3.2. Results

We estimate the preferences for each investor and report the distribution of estimated preferences across investors in Table 6. The mean preference parameters from our heterogeneous preferences specification are similar to the baseline estimates reported in Table 2. We also find persistent heterogeneity in preferences across investors, and we can reject the null hypothesis that investors have homogeneous preferences for each broker characteristic.

To help understand how preferences vary across investors, we project the parameters (α_i, β_i) on a vector of investor-specific characteristics D_i

$$\beta_i = \Gamma D_i + \eta. \tag{16}$$

The vector D_i captures the observable investor characteristics including whether the investor is a hedge fund, index fund, high churn/volume fund (above-average number of trades), high performing fund (above-average returns), or a large fund (above average size).[32],[33]

We report the corresponding estimates in Table 7. The estimates provide insight into how preferences vary across investors. For example, the preferences of hedge funds appear to be distinct from other institutional investors. Hedge fund managers do not appear to value sell-side research, placing little weight on the total number and the number of top research analysts a brokerage firm employs (columns 2 and 3). This result is intuitive, as hedge funds create value by conducting their own investment research and producing information in financial markets. Conversely, hedge funds appear to place greater value on informal research regarding order flow (column 5). While our previous results indicate that investors, on average, value brokers that are more central in the trading network, we find that instead hedge funds prefer to trade through less central brokers

[32]To identify index funds, we manually search the fund names in Ancerno for the word 'index' and flag the results with an indicator variable. Then, we aggregate this variable at the investment-company level by taking the average. Similarly, we identify hedge fund management companies in Ancerno using the procedure in Çötelioğlu et al. (2021). With the understanding that the identification is made at the management company-level, we label these firms "hedge funds" for short.

[33]We compute investors' six-month trading performance at the end of month t as the value-weighted return of all the trades executed over the prior six-month period evaluated at the end of the month in question. In particular, the percentage performance of all trades started by a manager over the prior six months is computed using closing prices at the end of month t, with sell trades' performance computed as the negative of a buy trade performance. We value-weight the performance of all the trades in the same six-month horizon ending in month t.

(column 4). One potential explanation for this finding is that hedge fund managers may be more concerned about concealing order flow and about brokers leaking their trades (Barbon, Di Maggio, Franzoni, and Landier (2019); and Di Maggio, Franzoni, Kermani, and Sommavilla (2019)). Thus, a hedge fund manager may prefer to trade through more peripheral brokers, conditioning on the broker possessing order flow information.

Index fund managers also have distinct preferences relative to other investors. Similar to hedge fund managers, index fund managers appear to place no value on sell-side research which is intuitive given that index fund managers have no use for research. We also find some evidence that index fund managers are among the most price-sensitive investors to both explicit and implicit trading costs (columns 1 and 6).

Overall, the results suggest that investor heterogeneity in brokerage markets is of first-order importance, especially when examining how investors value the ancillary services, such as sell-side research that brokerage firms offer. Accounting for this heterogeneity has important implications for how the proposed policies, such as the MiFID II regulations, will impact investors.

6.4. Other robustness

In Appendix A, we explore two additional robustness checks related to our definition of the market and whether the broker operates an alternative trading system (ATS).

First, we consider defining the trading market at a less granular level than at the sector-by-month level. One might think that investment managers make their trading decisions on a more aggregate level. In Appendix A.1 we re-estimate our baseline broker-choice specification where we define the market at the investor-by-year level, with the idea that investors make their execution decisions at a lower frequency than in our baseline analysis and across their entire portfolio. Overall, we find that changing the level of aggregation has little impact on our estimates (Table A1). The estimated preference parameters are quantitatively and qualitatively similar to our baseline estimates.

Second, we examine the impact of alternative trading systems (Appendix A.2). The central role of information in equity markets has helped lead to the proliferation of ATS and dark pools (see O'Hara and Ye (2011) and Zhu (2014)). Using data from the Securities and Exchange Commission, we track the development of ATSs in our sample. As of 2000, essentially none of the brokerage firms in our sample operated ATS; however, by 2014, roughly 50% of trades were executed through brokers who had access to their own ATS and dark pools. We re-estimate our baseline broker-choice specification where we control for whether a brokerage firm in our sample has access to its own ATS such as a dark pool. We construct the indicator variable ATS_{lt} which is equal to one if brokerage-firm l operates its own ATS at time t. An important caveat with our alternative trading system/dark pool variable is that it is measured at the broker-by-month level while our other brokerage firm characteristics are measured at the broker-by-month-by-sector level. Our demand estimates suggest that investors are willing to pay an additional 15% (2bps higher relative to the value of the transaction) fee per trade to trade through a brokerage firm that operates an

ATS (Table A2).

7. Soft dollars and management fees

Brokers traditionally provide bundled services to institutional investors by combining execution, research, and other brokerage services. Over the past 20 years, there has been a push among institutional investors and in policy circles to unbundle brokerage services to improve market competitiveness and transparency. Most recently, as part of MiFID II, European regulators are forcing brokers to unbundle their services. Bundling allows institutional investment managers to pay for research and other brokerage services with soft dollars through execution fees rather than to directly pay for these services with hard dollars. Soft-dollar transaction fees are not reported in the fund's expense ratio but are subtracted from the fund's returns.[34] The potential concern with soft-dollars payments is that they are borne by the end-investor and not disclosed by the fund. Hence, paying for research with soft dollars results in investment managers underreporting fund management fees. Here, we use our model from Section 2 and the corresponding estimates to quantify the value of soft dollars and determine how much investment managers are potentially underreporting fund management fees.

The term soft-dollar payments does not necessarily have a uniform definition in the industry and broadly incorporates two different types of research-related transactions (Blume (1993)). The first and most common type of transaction is when an investment manager uses broker commissions to pay a broker for research and other services that the broker produced in-house. In the second type of transaction, the investment manager uses broker commissions to pay for research and other services obtained from a third party. The broker then pays a portion of the corresponding commissions to the relevant third party. We use our framework to focus on soft-dollar payments for in-house research. We focus on these types of soft-dollar payments because they are more common (Blume (1993)) and can be more directly measured using our estimates.

The framework in Section 2 and the heterogeneous coefficient estimates from Section 6.3 (eq. 15) allow us to quantify soft-dollar in-house research-related payments in the brokerage industry. Our empirical estimates measure how each investment manager precisely values the in-house research produced by brokers, and how much more an investment manager is willing to pay on a per-transaction basis to have access to research. We then use these estimates to calculate how much larger the reported management fees would potentially be if investment managers were to include the value of soft-dollar in-house research-related payments in their management fees. Such analysis would not be possible without our structural model. For example, simply looking at the heterogeneity in fees (Figure 1) would be insufficient because we do not know if an investment manager pays a higher execution cost because the manager places a high value on research or because the manager is worse at execution. Our analysis allows us to precisely quantify soft-dollar research payments in terms of hard dollars.[35]

[34] http://www.finra.org/investors/funds-and-fees

[35] Note that in our unbundling counterfactual we focus on how much investors would be willing to pay for research

7.1. Quantifying soft dollars

We use the estimates from the demand model for brokerage services to quantify the total value investment managers obtain from having access to sell-side research. To calculate the total value of sell-side research, we compute the compensating variation required if we were to remove sell-side research from the marketplace altogether. The compensating variation tells us the amount institutional investors would be willing to pay in hard dollars to have access to sell-side research. We can then use the estimate of compensating variation to determine how much higher reported management fees would be if asset managers' clients were required to pay for research as part of the cost of investing in the fund. Alternatively, the estimates are informative about the gains that would accrue to clients if asset managers decided to shoulder the burden of paying for research out of their pocket.

Importantly, the compensating variation calculation is inherently a partial equilibrium calculation where the characteristics of brokers are held fixed. If regulators forced investors to pay for research with hard rather than soft dollars, the price of research in hard dollars in equilibrium would depend on competition among brokers and bargaining between investors and brokers, neither of which we have explicitly modeled. The advantage of focusing on compensating variation is that it can be directly calculated from our investor demand estimates without having to take a stance on the supply-side of the model or the nature of competition. To this end, compensating variation is informative of an investment manager's subjective value of research and provides an *upper bound* on how much management fees are currently underreported due to soft-dollar transactions.

We calculate the compensating variation at the investor-by-market level using our demand estimates. We calculate the compensating variation of investment manager i in sector k at time t as the expected profits of trading when the investment manager has access to sell-side research $(E[\pi_{ikt}])$ relative to the expected profits of trading when the investor does not have access to sell-side research $(E[\pi_{ikt}^{No\,Research}])$:

$$CV_{ikt}^{Research} = E[\pi_{ikt}] - E[\pi_{ikt}^{No\,Research}] \tag{17}$$

Following Petrin (2002), compensating variation in our discrete choice framework is given by

$$CV_{ikt}^{Research} = \frac{\ln\left(\sum_{l\in\mathcal{L}_{ikt}} exp(\overline{u}_{iklt})\right)}{\alpha_i} - \frac{\ln\left(\sum_{l\in\mathcal{L}_{ikt}} exp(\overline{u}_{iklt}^{No\,Resarch})\right)}{\alpha_i} \tag{18}$$

where $\overline{u}_{iklt} = -\alpha_i f_{iklt} + X_{klt}'\beta_i + \mu_{ilt} + \xi_{iklt}$ is the average utility derived by investment manager i from trading in sector k with broker l at time t and $\overline{u}_{iklt}^{No\,Resarch} = \overline{u}_{iklt} - X_{klt}^{Research}\beta_i^{Research} = -\alpha_i f_{iklt} + X_{klt}'\beta_i + \mu_{ilt} + \xi_{iklt} - X_{klt}^{Research}\beta_i^{Research}$ is the average utility derived by investment manager i from trading in sector k with broker l at time t excluding the utility from research

services in hard dollars. A related but separate interpretation of unbundling would be whether research services and transaction services are provided by the same broker. We focus on the former because it is more closely related to proposed regulations.

($X_{jkt}^{Research}\beta_i^{Research}$). Intuitively, the compensating variation is an increasing function of the utility of research ($X_{jkt}^{Research}\beta_i^{Research}$) aggregated across all brokers available to an investor in a given sector, $\mathcal{L}_{ikt}$. All else equal, the more utility an investment manager derives from research, the greater the required compensating variation. The scaling term $\frac{1}{\alpha_i}$ converts the required compensating utility in terms of profits/fees. Using the demand estimates (eq. 15), we calculate the compensating variation at the investor-by-market level.[36]

7.2. Results

Figure 5 plots the distribution of compensating variation at the investor-by-market level. For purposes of making an apples-to-apples comparison, we report the compensating variation for those markets where we observe at least one active research analyst. The average compensating variation is 3bps, which implies that the investment manager would be willing to pay an additional 3bps per trade to have access to sell-side research. Again, the value of research varies dramatically across the population of investment managers, with 25% of investors placing essentially no value (less than 0.5bps) on sell-side research. At the other extreme, 10% of investors would be willing to pay more than 7bps per trade to have access to outside research (Figure 5, Table 8).

We can use the compensating variation estimates to provide an estimate of the upper bound on how much higher management fees would be if investment managers had to pay for research with hard dollars and this cost was passed up to their clients. The compensating variation tells us investment managers' perceived value of the research that they consume through soft-dollar payments on a per-trade basis. Because our estimates of the value of research are on a per-trade basis, we annualize these implied research costs by multiplying them by the fraction of an investment manager's portfolio that is traded in a given year (the investment manager's portfolio turnover times two).[37],[38] Lastly, we compare the annualized implied research costs relative to the fund's annual management fees to determine how much investment managers underreport management

[36]Notice that in our demand specification we can write an investor's indirect utility as $\overline{u}_{iklt} = \ln(s_{iklt}) + \phi_{ikt}$, where ϕ_{ikt} is some market (investor-sector-time) specific constant. Thus we can compute the compensating variation empirically at the investor by market-level as

$$CV_{ikt}^{Research} = \left(\frac{\ln\left(\sum_{l\in\mathcal{L}_{ikt}} s_{ijkt}\right) - \ln\left(\sum_{l\in\mathcal{L}_{ikt}} s_{iklt} exp(-X_{klt}^{Research}\beta_i^{Research}\right)}{\alpha_i} \right)$$

where $X_{klt}^{Research}\beta_i^{Research}$ is the utility investor i derives from research.

[37]We calculate fund turnover and management fees for mutual funds as reported by CRSP Mutual Fund data. Because the Ancerno data is at the management-company level, but the mutual fund data is at the fund level, management companies in Ancerno (which we label investor) are matched to multiple mutual funds. We calculate the average turnover rate and manager expenses at the investor-by-year level where we take the equal weighted average across all of an investor's corresponding mutual funds. We calculate management fees for hedge funds as reported by TASS. We calculate portfolio turnover for hedge funds based on the average trading volume in our Ancerno sample.

[38]Fund turnover is calculated as the value of all transactions (buying, selling) divided by two, then divided by a fund's total holdings. Because we are interested in the number of trades that an investor makes in a given year, we multiply the investor's portfolio turnover by two to account for both sell (stocks removed from the portfolio) and buy trades (stocks added to the portfolio).

fees relative to the value they extract from soft-dollar research payments:

$$\frac{Annual\,Soft\,Dollars_{lt}}{Management\,Fees_{lt}} = \frac{\bar{CV}_{it}^{Research} \times Portfolio\,Turnover_{it} \times 2}{Management\,Fees_{t}}. \tag{19}$$

Figure 6 and Table 8 display our estimates of how much higher reported management fees would potentially be if investment managers had to explicitly pay for research payments in hard dollars instead of soft dollars and the cost was shifted onto their clients. Specifically, Figure 6 reflects the annual value of research obtained through soft-dollar payments relative to management fees at the investor-by-year level. The estimates indicate that reported management fees would be 4% higher if investment managers had to pay for the value of the research they consume with hard dollars. Again, there is substantial heterogeneity across investors. While there is minimal underreporting of management fees for 25% of our sample ($\frac{Annual\,Soft\,Dollars_{lt}}{Management\,Fees_{lt}} < 0.25\%$), management fees are underreported by more than 20% by some investment managers. For the investment managers in the top quartile (in terms of underreporting), reported management fees would be 15% higher if the funds had to pay for the research in hard dollars.

Because larger funds tend to place a higher value on research (Table 7), the results are even starker when we calculate the amount of underreporting weighted by assets under management (AUM), which may be the more relevant metric from an end investor's or policymaker's perspective. The third row of Table 8 displays the distribution of the value of soft dollars relative to management fees weighted by AUM. Overall, the results suggest that management fees would be 10% higher if investment managers had to pay for the value of the research they consume with hard dollars, and this cost was ultimately born by the fund's clients.

The evidence suggests that for many firms in our sample the value of soft-dollar research-related payments is substantial. Since the impetus behind MiFID II and its requirement for the unbundling of the services provided by brokers is to limit the use of soft dollars and improve market transparency, our results suggest that its effect might be significant in terms of how the overall cost of delegated asset management will change. Furthermore, one aspect emerging from our analysis that is often overlooked is that the effects of this regulation are likely to be disproportionate, as some funds are likely to be significantly more affected than others due to their tendency to compensate brokers for their research with trading commissions and are therefore more likely to be impacted.

8. Conclusion

Institutional investors continue to rely on high-touch brokerage transactions in equity markets even with the growth of alternative trading platforms. Given the sophistication of institutional investors and how well-developed equity markets are, why do institutional investors trade through brokers? This paper is a first step towards quantifying the value that brokers create for their clients.

Our results indicate that brokers create value for investors by providing efficient execution, sell-side research, and other informal services such as order flow information. Formal and informal research account for roughly half of the value that brokers create. While the average investor values these broker services, there is substantial heterogeneity across investors. Hedge funds place almost no value on sell-side research but place a large premium on order flow information. Conversely, large institutional investors are willing to pay up to 50% more per trade to access sell-side research.

Investors traditionally have paid for these research services with bundled commissions, or soft dollars, which potentially allow them to underreport their management fees. Our estimates suggest that investment management fees would be 10% higher if investment managers were forced to pay for the value of the research that they consume in hard rather than soft dollars. Overall, our results help explain why high-touch broker trading remains prominent in institutional equity markets.

References

Anand, A., Irvine, P., Puckett, A., Venkataraman, K., 2012. Performance of institutional trading desks: An analysis of persistence in trading costs. Review of Financial Studies 25 (2), 557–598.

Anand, A., Irvine, P., Puckett, A., Venkataraman, K., 2013. Institutional trading and stock resiliency: Evidence from the 2007–2009 financial crisis. Journal of Financial Economics 108 (3), 773–797.

Babus, A., Kondor, P., 2018. Trading and information diffusion in over-the-counter markets. Econometrica 86 (5), 1727–1769.

Barber, B., Lehavy, R., McNichols, M., Trueman, B., 2001. Can investors profit from the prophets? security analyst recommendations and stock returns. Journal of Finance 56 (2), 531–563.

Barber, B., Lehavy, R., McNichols, M., Trueman, B., 2003. Reassessing the returns to analysts' stock recommendations. Financial Analysts Journal 59 (2), 88–96.

Barbon, A., Di Maggio, M., Franzoni, F., Landier, A., 2019. Brokers and order flow leakage: Evidence from fire sales. Journal of Finance 74 (6), 2707–2749.

Benetton, M., 2021. Leverage regulation and market structure: An empirical model of the uk mortgage market. Journal of Finance forthcoming.

Berry, S., Levinsohn, J., Pakes, A., 1995. Automobile prices in market equilibrium. Econometrica, 841–890.

Berry, S. T., 1994. Estimating discrete-choice models of product differentiation. The RAND Journal of Economics, 242–262.

Bharath, S. T., Bonini, S., 2019. What is the value of analyst research? University of Arizona Working Paper.

Birru, J., Gokkaya, S., Liu, X., Stulz, R. M., 2019. Are analyst trade ideas valuable? Tech. rep., National Bureau of Economic Research.

Blume, M. E., 1993. Soft dollars and the brokerage industry. Financial Analysts Journal 49 (2), 36–44.

Boyarchenko, N., Lucca, D. O., Veldkamp, L., 2021. Taking orders and taking notes: Dealer information sharing in treasury auctions. Journal of Political Economy 129 (2), 607–645.

Buchak, G., Matvos, G., Piskorski, T., Seru, A., 2018. Fintech, regulatory arbitrage, and the rise of shadow banks. Journal of Financial Economics 130 (3), 453–483.

Chung, J.-W., Kang, B. U., 2016. Prime broker-level comovement in hedge fund returns: information or contagion? Review of Financial Studies 29 (12), 3321–3353.

Conrad, J. S., Johnson, K. M., Wahal, S., 2001. Institutional trading and soft dollars. Journal of Finance 56 (1), 397–416.

Çötelioğlu, E., Franzoni, F., Plazzi, A., 2021. What constrains liquidity provision? evidence from institutional trades. Review of Finance 25 (2), 485–517.

Coval, J. D., Moskowitz, T. J., 1999. Home bias at home: Local equity preference in domestic portfolios. Journal of Finance 54 (6), 2045–2073.

Cragg, J. G., Donald, S. G., 1993. Testing identifiability and specification in instrumental variable models. Econometric Theory 9 (2), 222–240.

Di Maggio, M., Franzoni, F., Kermani, A., Sommavilla, C., 2019. The relevance of broker networks for information diffusion in the stock market. Journal of Financial Economics 134 (2), 419–446.

Di Maggio, M., Kermani, A., Song, Z., 2017. The value of trading relations in turbulent times. Journal of Financial Economics 124 (2), 266–284.

Diamond, W., Jiang, Z., Ma, Y., 2020. A structural model of bank balance sheet synergies and the transmission of central bank policies. Available at SSRN.

Dick, A. A., 2008. Demand estimation and consumer welfare in the banking industry. Journal of Banking & Finance 32 (8), 1661–1676.

Duffie, D., Malamud, S., Manso, G., 2015. Information percolation in segmented markets. Journal of Economic Theory 158, 838–869.

Duffie, D., Manso, G., 2007. Information percolation in large markets. American Economic Review 97 (2), 203–209.

Edelen, R. M., Evans, R. B., Kadlec, G. B., 2012. Disclosure and agency conflict: Evidence from mutual fund commission bundling. Journal of Financial Economics 103 (2), 308–326.

Egan, M., 2019. Brokers vs. retail investors: Conflicting interests and dominated products. Journal of Finance 74 (3), 1217–1260.

Egan, M., Hortaçsu, A., Matvos, G., 2017a. Deposit competition and financial fragility: Evidence from the us banking sector. American Economic Review 107 (1), 169–216.

Egan, M., Lewellen, S., Sunderam, A., 2017b. The cross section of bank value. National Bureau of Economic Research WP 23921.

Egan, M., Matvos, G., Seru, A., 2019. The market for financial adviser misconduct. Journal of Political Economy 127 (1), 000–000.

Egan, M. L., Ge, S., Tang, J., 2020. Conflicting interests and the effect of fiduciary duty-evidence from variable annuities. Tech. rep., National Bureau of Economic Research.

Fang, B., Hope, O.-K., Huang, Z., Moldovan, R., 2020. The effects of MiFID II on sell-side analysts, buy-side analysts, and firms. Review of Accounting Studies 25 (3), 855–902.

Glode, V., Opp, C., 2016. Asymmetric information and intermediation chains. American Economic Review 106 (9), 2699–2721.

Goldstein, M. A., Irvine, P., Kandel, E., Wiener, Z., 2009. Brokerage commissions and institutional trading patterns. Review of Financial Studies 22 (12), 5175–5212.

Groysberg, B., Healy, P. M., 2013. Wall Street research: Past, present, and future. Stanford University Press.

Guo, Y., Mota, L., 2021. Should information be sold separately? evidence from MiFID II. Journal of Financial Economics forthcoming.

Hu, G., Jo, K. M., Wang, Y. A., Xie, J., 2018. Institutional trading and abel noser data. Journal of Corporate Finance 52, 143–167.

Irvine, P., Lipson, M., Puckett, A., 2006. Tipping. Review of Financial Studies 20 (3), 741–768.

Jegadeesh, N., Kim, J., Krische, S. D., Lee, C. M., 2004. Analyzing the analysts: When do recommendations add value? Journal of Finance 59 (3), 1083–1124.

Jiang, E. X., 2020. Financing competitors: Shadow banks' funding and mortgage market competition. Working Paper.

Koijen, R. S., Yogo, M., 2016. Shadow insurance. Econometrica 84 (3), 1265–1287.

Lang, M. H., Pinto, J., Sul, E., 2019. MiFID II unbundling and sell side analyst research. Available at SSRN 3408198.

Li, F. W., Mukherjee, A., Sen, R., 2017. Inside brokers HKUST Working Paper.

McNally, W. J., Shkilko, A., Smith, B. F., 2015. Do brokers of insiders tip other clients? Management Science 63 (2), 317–332.

Nanda, A., Groysberg, B., Prusiner, L., 2008. Lehman Brothers: Rise of the Equity Research Department. HBS.

Nevo, A., 2000. A practitioner's guide to estimation of random-coefficients logit models of demand. Journal of Economics & Management Strategy 9 (4), 513–548.

O'Hara, M., Ye, M., 2011. Is market fragmentation harming market quality? Journal of Financial Economics 100 (3), 459–474.

Perold, A. F., 1988. The implementation shortfall: Paper versus reality. Journal of Portfolio Management 14 (3), 4.

Petrin, A., 2002. Quantifying the benefits of new products: The case of the minivan. Journal of political Economy 110 (4), 705–729.

Puckett, A., Yan, X., 2011. The interim trading skills of institutional investors. Journal of Finance 66 (2), 601–633.

Robles-Garcia, C., 2019. Competition and incentives in mortgage markets: The role of brokers. London School of Economics Working paper.

Stickel, S. E., 1992. Reputation and performance among security analysts. Journal of Finance 47 (5), 1811–1836.

Stock, J. H., Yogo, M., 2005. Testing for weak instruments in linear IV regression. Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg, 80.

Wagner, W. H., Edwards, M., 1993. Best execution. Financial Analysts Journal 49 (1), 65–71.

Wang, Y., Whited, T. M., Wu, Y., Xiao, K., 2020. Bank market power and monetary policy transmission: Evidence from a structural estimation. NBER Working Paper 27258.

Womack, K. L., 1996. Do brokerage analysts' recommendations have investment value? Journal of Finance 51 (1), 137–167.

Xiao, K., 2020. Monetary transmission through shadow banks. Review of Financial Studies 33 (6), 2379–2420.

Xiao, K., Whited, T., Wu, Y., 2020. Risk incentives for banks with market power when interest rates are low. Working Paper.

Ye, M., Zhu, W., 2019. Where do activist shareholders trade? University of Illinois Working Paper.

Zhu, H., 2014. Do dark pools harm price discovery? Review of Financial Studies 27 (3), 747–789.

Figure 1: Brokerage Fees

(a) Fees (bp of Transaction Value)

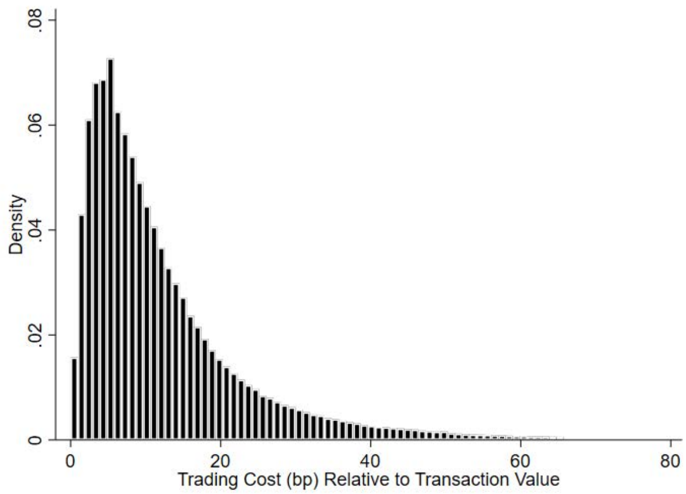

(b) Fees ($ per Share)

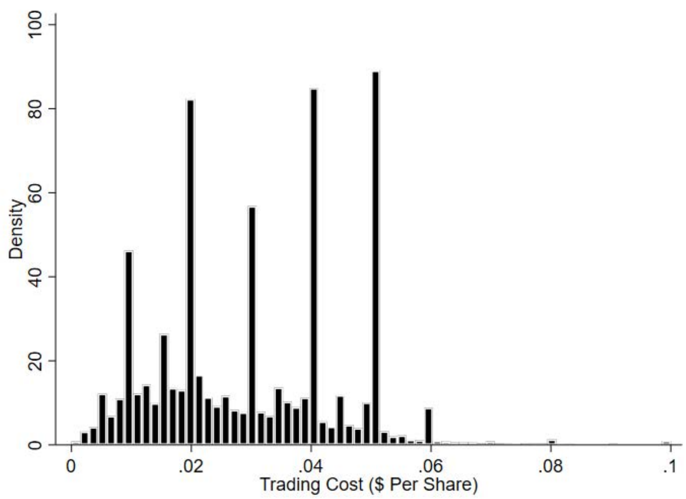

Note: Figure 1 displays the distribution of fees charged by brokerage firms in terms of the cost relative to the value of the transaction and the cost in terms of dollars per share. Observations are averaged at the investor-by-broker-by-sector-by-month level which is the unit of observation in our main analysis.

Figure 2: Distribution of Broker Fixed Effects

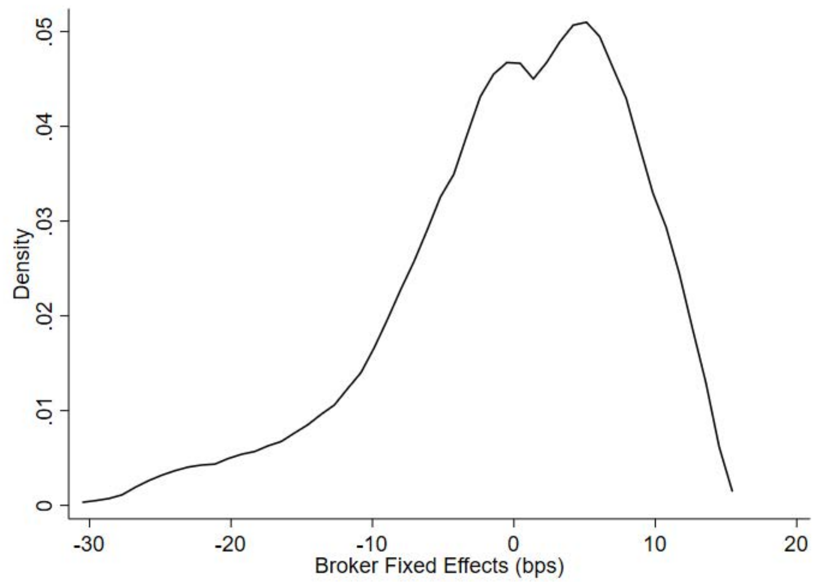

Note: Figure 2 displays the distribution of estimated broker fixed effects corresponding to column (3) in Table 2. The fixed effects are scaled by $1/\alpha$, which puts the fixed effects in terms of basis points (bps) rather than utils. Observations are averaged at the investor-by-broker-by-sector-by-month level which is the unit of observation in our main analysis.

Figure 3: Value Decomposition

(a) Value of Research and Execution

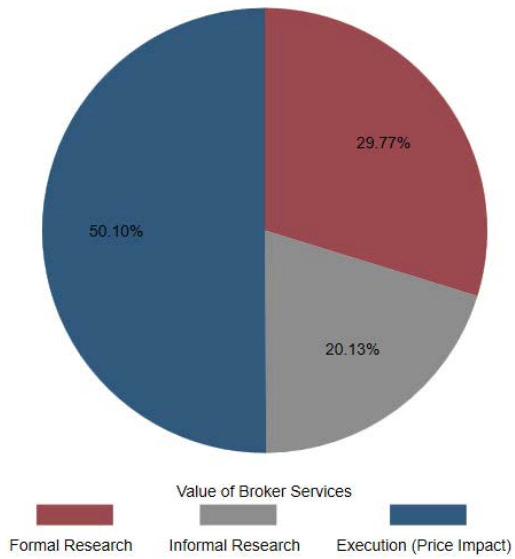

(b) Value of Research, Execution, and Broker Fixed Effects

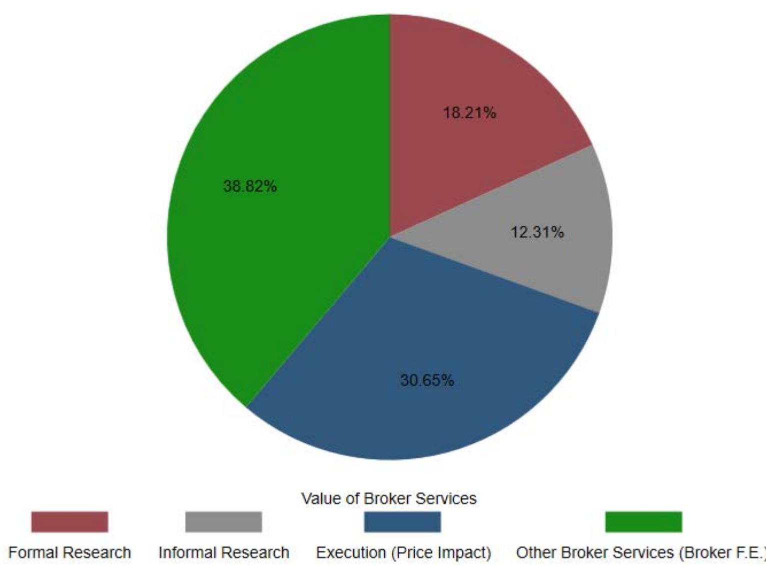

Note: Figure presents the average utility flow generated by broker services per transaction. For each transaction, we calculate the utility flow generated by broker services as $X'\hat{\beta}$, which is measured in utils. To aid in interpretation we scale utility flow generated by each service relative to the total utility flow. In panel(a) we calculate the value of formal research as the utility flow generated by analysts and top analysts; the value of informal research as the utility flow generated by trading with informed brokers and more central brokers; and the value of execution based on our measure of Price Impact. In panel (b) we calculate the value of formal research as the utility flow generated by analysts and top analysts; the value of informal research as the utility flow generated by trading with informed brokers and more central brokers; the value of execution based on our measure of Price Impact; and the value of other broker services based on our estimated broker-time fixed effects ($\hat{\mu_{lt}}$). The estimates $\hat{\beta}$, $\hat{\alpha}$, and $\hat{\mu_{lt}}$ correspond to the estimates reported in column (4) of Table 2.

Figure 4: Elasticity of Demand over Time

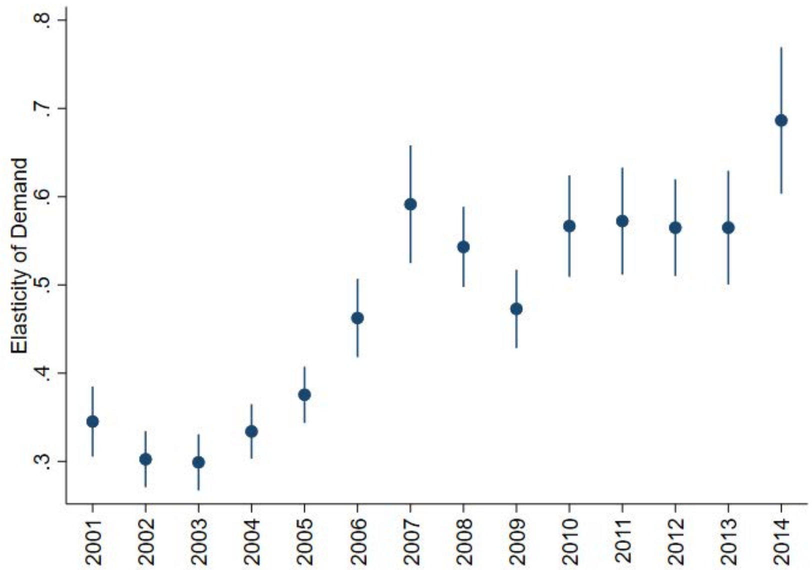

Note:Figure 4 presents the elasticity of demand over time. We compute the elasticity of demand as $\alpha_t(1-\bar{s})\bar{f}$, where $\bar{s}$ and $\bar{f}$ correspond to the average market share and fee in the sample. The estimates of α_t correspond to the estimates reported in Table 5. Confidence intervals are computed at the 95% level and standard errors are clustered at the broker-by-year level.

Figure 5: Total Value of Research

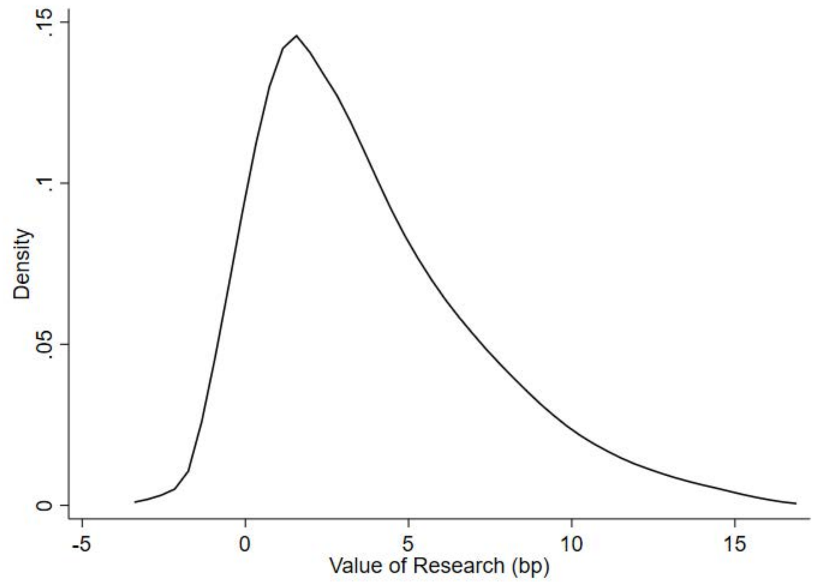

Note: Figure 5 presents the distribution of compensating variation if we were to remove sell-side research from the market. The compensating variation indicates how much each investor would need to be compensated on a per-trade basis to make them indifferent between a regime with and without sell-side research. We compute the compensating required for each investor at the market-level according to eq. (18). Observations are at the investor-by-month-by-sector level. The above figure displays the distribution truncated at the 1% and 99% level.

Figure 6: Research Related Soft Dollars Relative to Management Fees

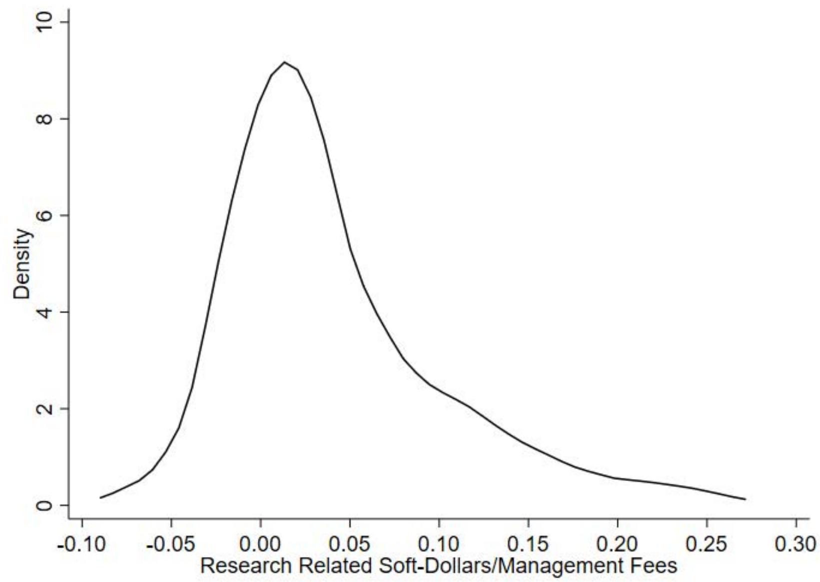

Note: Figure 6 presents the distribution of the annual value of soft-dollar research payments relative to the investor's management fees. Observations are at the investor-by-year level. We calculate the annual value of soft-dollar research payments based on the compensating variation required if we were to remove sell-side research from the market (eq. 18; Table 5). Specifically, we calculate the annual value of soft-dollar research related payments as the average compensating variation at the investor-by-year level multiplied by how often the institutional investor turns over his/her portfolio. The above figure displays the distribution truncated at the 2.5% and 97.5% level.

Table 1: Summary Statistics

Variable	Obs	Mean	Std. Dev.
Fees($ per share)	6,861,419	0.030	0.016
Fees (%)	6,861,419	0.13%	0.13%
Price Impact	6,861,419	0.19%	0.64%
Access to an ATS/Dark Pool	6,861,419	25.59%	43.63%
Research Analysts:			
Number of Analysts	6,861,419	1.48	2.39
Number of Top Analysts	6,861,419	0.17	0.48
Broker Information:			
Eigenvector Centrality	6,861,419	0.052	0.10
Informed Broker (Di Maggio et al. 2018)	6,861,419	28%	45%
Equity Traders:			
Number of Traders	3,377,309	255	238
Pct of Traders Receiving Misconduct Disclosures	3,397,871	0.20%	0.61%
Average Trader Experience	3,377,309	11.65	2.66
Distance (miles)	2,010,963	669	804
Close Distance (Dist.<100 miles)	2,010,963	33%	47%
Institutional Investors:			
Hedge Fund	6,861,419	0.21	0.41
Index Fund	6,861,419	0.028	0.10
Number of Trading Partners (Per Market)	6,861,419	17.06	11.94

Note: Table 1 displays the summary statistics corresponding to our data set. Each variable is described in detail in Section 4.2. Observations are at the investor-by-month-by-sector-by-broker level.

Table 2: Broker Choice

	(1)	(2)	(3)	(4)	(5)
Fees	-152***	-413***	-402***	-403***	-374***
	(4.46)	(7.95)	(7.14)	(7.07)	(5.44)
Price Impact:	3.86***	-38.8*	-15.6	-21.0*	-20.0
	(0.38)	(20.8)	(13.4)	(12.6)	(13.0)
Research:					
Number of Analysts	0.067***	0.070***	0.031***	0.035***	0.032***
	(0.0038)	(0.0038)	(0.0025)	(0.0018)	(0.0016)
Number of Top Rated Analysts	0.15***	0.15***	0.066***	0.065***	0.075***
	(0.0099)	(0.011)	(0.0059)	(0.0043)	(0.0044)
Information:					
Eigenvector Centrality	1.27***	1.16***	0.50***	0.30***	0.63***
	(0.066)	(0.068)	(0.040)	(0.045)	(0.036)
Informed Broker	0.31***	0.26***	0.12***	0.10***	0.098***
	(0.015)	(0.015)	(0.0066)	(0.0042)	(0.0041)
Sector×Investor×Time Fixed Effects	X	X	X	X	X
Broker Fixed Effects			X		
Broker×Time Fixed Effects				X	
Broker×Investor×Time Fixed Effects					X
IV (Commissons & Price Impact)		X	X	X	X
Cragg Donald F-Statistic for IV		6,205	2,881	1,297	1,166
Observations	6,187,028	5,607,856	5,607,852	5,607,323	5,487,670
R-squared	0.302	0.270	0.299	0.315	0.587
Mean Elasticity with Respect to Fees	0.18	0.48	0.47	0.47	0.44
Value of 1σ Decrease in Price Impact (bp)	-1.62	6.01	2.48	3.33	3.42
Value of Research:					
Value of an Additional Analyst (bp)	4.41	1.69	0.77	0.87	0.86
Value of an Additional Top Analyst (bp)	14.28	5.33	2.41	2.48	2.86
Value of Information:					
Value of 1σ Increase in Eigenvector Centrality (bp)	8.33	2.80	1.24	0.74	1.68
Value of an Informed Broker (bp)	20.39	6.30	2.99	2.48	2.62

Note: Table 2 displays the estimation results corresponding to our discrete choice broker model (eq. 8). The unit of observation is at the investment manager-by-broker-by-month-by-sector (6-digit GICS) level over the period 2001-2014. Each independent variable is described in detail in Section 4.2. We measure fees in percentage terms relative to the value of the transaction. As described in the text, we instrument for fees using the average historical fee charged by the broker in terms of cents per share divided by the share price of the stock being traded. We instrument for price impact using the lagged price impact. Standard errors are clustered at the broker-by-year level and are reported in parentheses. The symbols *,**, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.
In the bottom half of the table, we interpret the coefficient estimates. Elasticity of demand is calculated as the average of $-\alpha * (1 - s) * fee$. We compute the value of each independent variable as the ratio of the coefficient of interest divided by an investor's sensitivity with respect to fees ($-\alpha$). For example, we calculate the value of an analyst in column (1) as 10,000 × 0.067/152 = 4.41bps.

Table 3: Broker Choice and Trader Characteristics

	(1)	(2)	(3)	(4)
Fees	-481***	-482***	-482***	-482***
	(10.3)	(10.3)	(10.3)	(10.3)
Trader Characteristics:				
Misconduct	-2.18**			-1.83*
	(1.00)			(1.01)
Trader Experience		0.22***		0.21***
		(0.040)		(0.042)
Trader Experience^2		-0.0080***		-0.0077***
		(0.0015)		(0.0015)
Number of Traders (100s)			-0.0054	0.055
			(0.042)	(0.050)
Number of Traders$^2 (100s)$			-0.0033	-0.0074*
			(0.0035)	(0.0041)
Sector×Investor×Time Fixed Effects	X	X	X	X
Other Controls	X	X	X	X
Broker Fixed Effects	X	X	X	X
IV (Commissons & Price Impact)	X	X	X	X
Cragg Donald F-Statistic for IV	1,105	1,017	1,078	1,003
Observations	3,134,050	3,120,165	3,134,050	3,120,165
R-squared	0.294	0.293	0.294	0.293
Mean Elasticity	0.56	0.56	0.56	0.56
Value of Trader Characteristics:				
1pp Inc. in Misc. (bp).	-0.45			-0.38
1 Year Inc. in Trader Experience (bp):		0.70		0.97
100 Inc. in Number of Traders			-0.45	0.37

Note: Table 3 displays the estimation results corresponding to our discrete choice broker model (eq. 8). The unit of observation is at the investment manager-by-broker-by-month-by-sector level over the period 2001-2014. The independent variable Misconduct measures the share of equity traders working for the brokerage firm in a given year that receive misconduct disclosures. Trader Experience measures the average trader experience in years. Number of Traders measures the number of traders working at a brokerage firm and is measured in 100s of traders. Other controls include: Price Impact, Number of Research Analysts, Number of Top Research Analysts, Broker Eigenvector Centrality, and Informed. We instrument for fees and price impact as described in the text. Standard errors are clustered at the broker-by-year level and are reported in parentheses. The symbols *,**, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.
In the bottom half of the table, we interpret the coefficient estimates. Elasticity of demand is calculated as the average of $-\alpha * (1-s) * fee$. We compute the value of each independent variable as the ratio of the coefficient of interest divided by an investor's sensitivity with respect to fees (-α). We calculate the marginal value of a year of Trader Experience at the average value of Trader Experience (11.65 years). Similarly, we calculate the marginal value of an additional 100 traders at the average value of Number of Traders (250 traders).

Table 4: Broker Choice and Distance

	(1)	(2)	(3)	(4)
Fee	-148***	-404***	-400***	-398***
	(7.72)	(11.5)	(9.92)	(9.65)
Close Distance (Less than 100 miles)	0.41***	0.42***	0.34***	0.35***
	(0.048)	(0.051)	(0.053)	(0.054)
Sector×Investor×Time Fixed Effects	X	X	X	X
Other Controls	X	X	X	X
Broker Fixed Effects			X	
Broker ×Time Fixed Effects				X
IV (Commissons & Price Impact)		X	X	X
Cragg Donald F-Statistic for IV		1,493	743	305
Observations	1,910,011	1,816,638	1,816,637	1,816,374
R-squared	0.299	0.281	0.308	0.340
Mean Elasticity	0.17	0.47	0.47	0.46
Value Being Less than 100 miles (bp)	27.70	10.40	8.50	8.79

Note: Table 4 displays the estimation results corresponding to our discrete choice broker model (eq. 8). The unit of observation is at the investment manager-by-broker-by-month-by-sector level over the period 2001-2014. Close Distance is a dummy variable indicating that the broker and investor are located within 100 miles of each other. Other controls include: Price Impact, Number of Research Analysts, Number of Top Research Analysts, Broker Eigenvector Centrality, and Informed. We instrument for fees and price impact as described in the text. Standard errors are clustered at the broker-by-year level and are reported in parentheses. The symbols *,**, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.
In the bottom half of the table, we interpret the coefficient estimates. Elasticity of demand is calculated as the average of $-\alpha * (1-s) * fee$. We compute the value of Distance as the ratio of the Distance coefficient divided by an investor's sensitivity with respect to fees ($-\alpha$).

Table 5: Broker Choice: Time-varying Preference Parameters

Indep Var.	Fees	Price Impact	Number of Analysts	Number of Top Analysts	Eigenvector Centrality	Informed Broker
Year:						
2001	-297***	-48.9	0.055***	0.0098	0.71	0.13***
	(17.4)	(40.4)	(0.016)	(0.0081)	(0.48)	(0.022)
2002	-260***	-69.3**	0.049***	0.029***	0.54**	0.12***
	(13.9)	(29.7)	(0.015)	(0.0082)	(0.24)	(0.022)
2003	-257***	-368	0.086***	0.039***	0.32	0.16***
	(14.0)	(239)	(0.019)	(0.0083)	(0.20)	(0.037)
2004	-287***	-73.9	0.077***	0.036***	0.24**	0.13***
	(13.6)	(78.1)	(0.018)	(0.0075)	(0.12)	(0.017)
2005	-323***	0.57	0.083***	0.045***	0.65***	0.14***
	(14.0)	(48.0)	(0.0091)	(0.0085)	(0.10)	(0.015)
2006	-398***	-91.5*	0.045***	0.048***	0.31***	0.12***
	(19.5)	(52.9)	(0.0097)	(0.0066)	(0.12)	(0.016)
2007	-509***	-19.9	0.045***	0.038***	0.031	0.079***
	(29.3)	(30.5)	(0.015)	(0.0051)	(0.18)	(0.014)
2008	-467***	63.3	0.055***	0.032***	-0.075	0.065***
	(19.9)	(46.4)	(0.016)	(0.0053)	(0.12)	(0.014)
2009	-407***	132*	0.053**	0.041***	-0.15	0.093***
	(19.5)	(67.7)	(0.023)	(0.0063)	(0.18)	(0.017)
2010	-488***	122*	0.072***	0.029***	0.11	0.071***
	(25.3)	(68.2)	(0.017)	(0.0071)	(0.091)	(0.012)
2011	-493***	51.1	0.061***	0.033***	0.0068	0.087***
	(26.6)	(44.9)	(0.016)	(0.0054)	(0.13)	(0.014)
2012	-486***	27.7	0.063***	0.029***	0.31***	0.10***
	(24.1)	(28.5)	(0.018)	(0.0053)	(0.11)	(0.015)
2013	-486***	-57.3	0.063***	0.044***	0.21	0.085***
	(28.3)	(53.1)	(0.018)	(0.0076)	(0.16)	(0.015)
2014	-591***	16.2	0.078***	0.030***	0.21*	0.10***
	(36.4)	(58.3)	(0.020)	(0.0050)	(0.11)	(0.017)
Sector×Investor×Time Fixed Effects	X					
Broker×Time Fixed Effects	X					
IV (Commissons & Price Impact)	X					
Observations	5,607,323					
R-squared	0.254					

Note: Table 5 displays the estimation results corresponding to our discrete choice broker model (eq. 8), where we allow investor preferences to vary from year-to-year. The table presents the results for a single model/regression. Each cell of the table corresponds to a coeffıcent from the regression. The unit of observation in the regression is at the investment manager-by-broker-by-month-by-sector (6-digit GICS) level over the period 2001-2014. Each independent variable is described in detail in Section 4.2. We measure fees in percentage terms relative to the value of the transaction. As described in the text, we instrument for fees using the average historical fee charged by the broker in terms of cents per share divided by the share price of the stock being traded. We instrument for price impact using the lagged price impact to account for measurement error. Standard errors are clustered at the broker-by-year level and are reported in parentheses. The symbols *,**, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 6: Broker Choice - Heterogeneous Coefficients

	Mean	Std. Dev.	F Stat.
Fees	-435.67***	428.14***	73.1
Price Impact:	11.35	309.64***	2.02
Research			
Number of Analysts	0.011	0.040***	13.6
Number of Top Rated Analysts	0.038***	0.089***	5.34
Information:			
Eigenvector Centrality	0.51***	0.83***	12.7
Informed Broker	0.08	0.13***	7.7
Sector×Investor×Time Fixed Effects	X		
Broker×Investor Fixed Effects	X		
IV (Commissons)	X		
Observations	5,560,530		

Note: Table 6 displays the estimation results corresponding to our heterogeneous coefficient discrete choice broker model (eq. 15). The unit of observation is at the investor-by-broker-by-month-by-sector (6-digit GICS) level over the period 2001-2014. We restrict our analysis to the 243 investors where we observe at least 1,000 observations. Here, we allow preferences to vary across investors. Consequently, we report the mean and standard deviation of preferences across the investors in our sample. To control for outliers, we report the estimated coefficients winsorized at the 1% level. Each independent variable is described in detail in Section 4.2. We measure fees in percentage terms relative to the value of the transaction. For each broker characteristic, we report the F Statistic corresponding to the null hypothesis that all investors have the same preferences over the given broker characteristics. The symbols *,**, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 7: Investor Preferences by Type of Investor

	Fees (1)	Analyst (2)	Top Analyst (3)	Centrality (4)	Informed (5)	Price Impact (6)
Hedge Fund	186*	-0.0062	-0.043***	-0.42***	0.089***	67.9
	(110)	(0.0066)	(0.015)	(0.12)	(0.034)	(49.6)
Index Fund	-109	-0.0056	-0.093*	0.44	-0.062	-168
	(257)	(0.021)	(0.055)	(0.75)	(0.062)	(164)
Large Investor	76.4	0.015**	0.037***	0.12	0.027**	-21.4
	(93.9)	(0.0064)	(0.013)	(0.088)	(0.013)	(27.0)
High Performance	36.1*	-0.000089	-0.0042	0.13**	0.00049	5.90
	(21.6)	(0.0035)	(0.0068)	(0.059)	(0.0063)	(13.5)
High Churn	36.9	0.0084*	0.034***	0.053	0.016	24.0
	(91.3)	(0.0049)	(0.012)	(0.078)	(0.010)	(23.5)
Constant	-620***	0.015**	0.040***	0.49***	0.071***	-8.11
	(222)	(0.0063)	(0.014)	(0.11)	(0.015)	(27.0)
Observations	243	243	243	243	243	243
R-squared	0.023	0.064	0.178	0.110	0.146	0.040

Note: Table 7 presents the results corresponding to a linear regression (eq. 16) where we examine how investor preferences vary with observable investor characteristics. Observations are at the investor-level and the dependent variable in each column corresponds to the investors' preferences β_i for a given broker characteristic. The estimates of investor preferences correspond to the results reported in Table 6. Because the dependent variable is estimated from the data, we weight observation based on the number of observations we have for each investor. To account for outliers, we winsorize the estimated parameters at the 1% level. The independent variables Hedge Fund, High Churn (above average number of trades), High Performance (above average returns), and Large Investor (above average size) are all dummy variables. The variable Index Fund is between zero and one and indicates whether the investor operates one or more index funds. Specifically, we calculate index manually by searching the fund names in Ancerno for the word 'index' and flag the results with an indicator variable. Then, we aggregate this variable at the investment company-level by taking the average. The symbols *,**, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 8: Research Related Soft Dollars Relative to Management Fees

	Mean	SD	Percentile				
			10th	25th	50th	75th	90th
Soft-Dollars:							
Per Trade (bp)	2.80	4.35	0.07	0.61	1.92	4.09	7.13
Annual (% of Management Fees)	4.33	7.43	-1.97	0.15	2.20	7.06	14.28
Annual (% of Management Fees) Weighted by AUM	9.79	9.06	0.07	2.28	7.82	15.07	27.10

Note: Table 8 presents the distribution of the value of soft-dollar research payments on a per-trade basis (in bp) and annualized (% of management fees) for the investment managers in our sample. Observations for soft-dollars per trade are at the investor-by-month-by-sector level, matching the unit of observation corresponding to our estimates reported in Tables 2-6. We calculate the value of soft-dollar research payments on a per-trade basis based on the compensating variation required if we were remove sell-side research from the market (eq. 18; Table 5). Observations for annual soft-dollars are at the investor-by-year level. We calculate the annual value of soft-dollar research related payments as the average compensating variation at the investor-by-year level multiplied by how often the institutional investor turns over his/her portfolio. We express the annual value of soft-dollar research payments relative to annual management expenses. To account for outliers, we winsorize annual soft-dollars at the 2.5% level. In the final row, we calculate the distribution of the annual value of soft-dollar payments weighted by the investment manager's assets under management.

A. Additional Robustness

A.1. Alternative Market Share Definitions

As a robustness check, we consider defining the trading market at a less granular level than at the sector-by-month level. In Table A1 we re-estimate our baseline broker-choice specification where we define the market at the investor-by-year level, with the idea that investors make their execution decisions at a lower frequency than in our baseline analysis and across their entire portfolio. Overall, we find that changing the level of aggregation has little impact on our estimates.

A.2. Alternative Trading Systems

We also examine how the proliferation of alternative trading systems (ATS) impacts investors' trading behavior. Using data from the Securities and Exchange Commission, we track the development of ATSs in our sample. We construct the indicator variable ATS_{lt} which is equal to one if brokerage-firm l operates its own ATS at time t. While few brokers operated ATS as of 2000, by 2014, roughly 50% of trades were executed through brokers who had access to their own ATS and dark pools (Figure A1). As a robustness check, we re-estimate our demand specification where we include the control variable ATS_{lt} and report the corresponding estimates in Table (Table A2). The estimates in column (2) suggest that investors are willing to pay an additional 15% (2bps higher relative to the value of the transaction) fee per trade to trade through a brokerage firm that operates an ATS.

A.3. Instruments for Research Coverage

As discussed in Section 4.2, a natural concern in the demand estimation literature is the endogeneity of prices or, in this setting, fees. However, one may also be concerned that other broker characteristics, such as research coverage could also be endogenous. As a robustness check, we include additional specifications where we instrument for research coverage using historical coverage. Specifically, we instrument for the number of analysts and number of top analysts a broker employs in a given sector and month using the 12-month lagged values. Conceptually, we would like a broker-sector-specific instrument that is correlated with the cost of providing analyst coverage but that is orthogonal to demand shocks. The rationale behind using lagged research coverage as an instrument for current research coverage is that it is likely to be relevant because it's costly to adjust research coverage and the costs of providing research analyst coverage are correlated over time, and it is potentially exogenous because lagged research coverage is likely to be uncorrelated with current demand shocks. The concern would be that if demand shocks are highly serially correlated and research coverage is endogenous, then lagged research coverage could be correlated with current demand shocks.

We report the corresponding instrumental variables specifications in Table A4. Consistent with our baseline estimates, we find that investors are more likely to trade through brokers that have analyst coverage and those that employ the top analysts. The magnitudes are quite similar to our baseline estimates. For example, in our most stringent specification with investor-broker-time fixed effects the instrumental variable results indicate that investors behave as if they are willing to pay an additional 1.12bps to have access to an analyst and 3.00bps to have access to a top analyst (Table A4, Column 4). Conversely, our baseline estimates, where we do not instrument for research coverage, indicate that investors behave as if they are willing to pay 0.86bps to have access to an analyst and 2.86bps to have access to a top analyst (Table 2, Column 5).

A.4. Large Price Impact

As an additional robustness check, we examine an investors' sensitivity to large trading costs. Specifically, we re-estimate our baseline specification where we control for implicit trading costs with the variable $LargePriceImpact_{lkt}$, which measures whether the price impact was greater than 0.25% (roughly the 50th percentile). We report the estimates in Table A5. The results in column (2) suggest that investors behave as though they are willing to pay an additional 34bps (=1.35/400) higher broker fee, relative to the value of the transaction, to avoid a large price impact of at least 0.25%. Given that the median (mean) large price impact is 0.46% (0.66%), this implies that investors trade off implicit and explicit trading costs almost one-for-one. The results suggest that investors avoid brokers with a track record of particularly poor execution, which may be more salient and predictable than average execution.

B. Accounting for the Outside Good

In the baseline setup of the model, we assume that an investor has a trade she needs to execute, and she can route the order through any of the n available brokers. In other words, the investor has decided to trade, and we are estimating her execution decision conditional on her initial decision to trade. The investor's indirect utility of executing trade j in industry sector k through brokerage firm l at time t is given by:

$$E[u_{ijklt}] = -\alpha_i f_{ijklt} + X'_{klt}\beta_i + \mu_{ilt} + \xi_{iklt} + \epsilon_{ijklt}.$$

Under the assumption that the investor-trade specific utility shock ϵ_{ijklt} is distributed T1EV, the probability that an investor executes trade j through broker l is given by:

$$P(l) = \frac{\exp\left(-\alpha_i f_{iklt} + X'_{klt}\beta_i + \mu_{ilt} + \xi_{iklt}\right)}{\sum_{m\in\mathcal{L}} \exp\left(-\alpha_i f_{ikmt} + X'_{kmt}\beta_i + \mu_{imt} + \xi_{ikmt}\right)}.$$

We can then write the share of trades executed through broker l as in logs as

$$\ln s_{iklt} = -\alpha_i f_{iklt} + X'_{klt}\beta_i + \mu_{ilt} + \xi_{iklt} - \ln\left(\sum_{m\in\mathcal{L}} \exp\left(-\alpha_i f_{ikmt} + X'_{kmt}\beta_i + \mu_{imt} + \xi_{ikmt}\right)\right)$$

Notice that the non-linear term $\ln\left(\sum_{m\in\mathcal{L}} \exp\left(-\alpha_i f_{ikmt} + X'_{kmt}\beta_i + \mu_{imt} + \xi_{ikmt}\right)\right)$ is constant in a given market. Thus, we can estimate the above equation as

$$\ln s_{iklt} = -\alpha_i f_{iklt} + X'_{klt}\beta_i + \mu_{ilt} + \mu_{ikt} + \xi_{iklt}.$$

Where μ_{ikt} is an investor-by-market fixed effect, which absorbs the nonlinear term. Thus, we can recover the preference parameters α and β by estimating the model via OLS/IV.

While we cast our model in the context of an investor's decision regarding where to execute her trade *conditional* on the initial decision to trade a specific security, the model and corresponding estimates also generalize to the setting where brokers influence an investor's initial decision of whether or not to trade. Furthermore, because we include investor-market fixed effects in our estimation procedure, adding the outside option of not trading would produce numerically equivalent estimation results.

To see this, consider the alternative formulation of the investor's problem where investors now have the outside option of not trading, which yields utility $u_{ik0t.}$. The probability that an investor trades through brokerage firm l, denoted $\widetilde{P}(l)$, is then

$$\widetilde{P}(l) = \frac{\exp\left(-\alpha_i f_{iklt} + X'_{klt}\beta_i + \mu_{ilt} + \xi_{iklt}\right)}{exp(u_{ik0t}) + \sum_{m\in\mathcal{L}} \exp\left(-\alpha_i f_{ikmt} + X'_{kmt}\beta_i + \mu_{imt} + \xi_{ikmt}\right)}.$$

Notice that the difference between $\widetilde{P}(l)$ and $P(l)$ is the denominator. $\widetilde{P}(l)$ includes the term $exp(u_{ik0t})$ in the denominator, which reflects the utility an investor gets if she does not trade.

Let $\widetilde{s}_{iklt}$ denote the share of trades executed through brokerage firm l, which corresponds to $\widetilde{P}(l)$, where we have computed the share of trades accounting for the outside good option (i.e., not trading). We can again then write the share of trades executed through broker l as in logs as

$$\begin{aligned}\ln \widetilde{s}_{iklt} =& -\alpha_i f_{iklt} + X'_{klt}\beta_i + \mu_{ilt} + \xi_{iklt} \\ & - \ln\left(exp(u_{ik0t}) + \sum_{m\in\mathcal{L}} \exp\left(-\alpha_i f_{ikmt} + X'_{kmt}\beta_i + \mu_{imt} + \xi_{ikmt}\right)\right).\end{aligned}$$

Note that mechanically it is the case that $\widetilde{s}_{iklt} = s_{iklt} \times (1 - \widetilde{s}_{ik0t})$, where we have scaled the conditional share (s_{iklt}) by the probability the investor trades $(1 - \widetilde{s}_{ik0t})$. Thus, the dependent variable can be written as $\widetilde{s}_{iklt} = s_{iklt} \times (1 - \widetilde{s}_{ik0t})$ such that we have

$$\begin{aligned}\ln s_{iklt} + \ln(1 - \widetilde{s}_{iklt}) =& -\alpha_i f_{iklt} + X'_{klt}\beta_i + \mu_{ilt} + \xi_{iklt} \\ & - \ln\left(exp(u_{ik0t}) + \sum_{m\in\mathcal{L}} \exp\left(-\alpha_i f_{ikmt} + X'_{kmt}\beta_i + \mu_{imt} + \xi_{ikmt}\right)\right).\end{aligned}$$

Both the term $\ln\left(exp(u_{ik0t}) + \sum_{m\in\mathcal{L}} \exp\left(-\alpha_i f_{ikmt} + X'_{kmt}\beta_i + \mu_{imt} + \xi_{ikmt}\right)\right)$ on the RHS of the above equation and the term $\ln(1 - \widetilde{s}_{iklt})$ on the LHS of the above equation are constant within investor-market (ikt triplet). Consequently, if we include an investor-market fixed effect (ϕ_{ikt}) when estimating the above equation we get:

$$\ln \widetilde{s}_{iklt} = -\alpha_i f_{iklt} + X'_{klt}\beta_i + \mu_{ilt} + \phi_{ikt} + \xi_{iklt},$$

which produces numerically equivalent results as estimating:

$$\ln s_{iklt} = -\alpha_i f_{iklt} + X'_{klt}\beta_i + \mu_{ilt} + \mu_{ikt} + \xi_{iklt}.$$

where $\mu_{ikt} = \phi_{ikt} - \ln(1 - \widetilde{s}_{ik0t})$. Consequently, if we were to extend our model and data to allow for the option of not trading, our estimation procedure would produce numerically equivalent estimates of α, β, and μ_{ilt} regardless of how we define the outside option.

C. Additional Figures and Tables

Figure A1: Share of Brokers that Operate Dark Pools Over Time

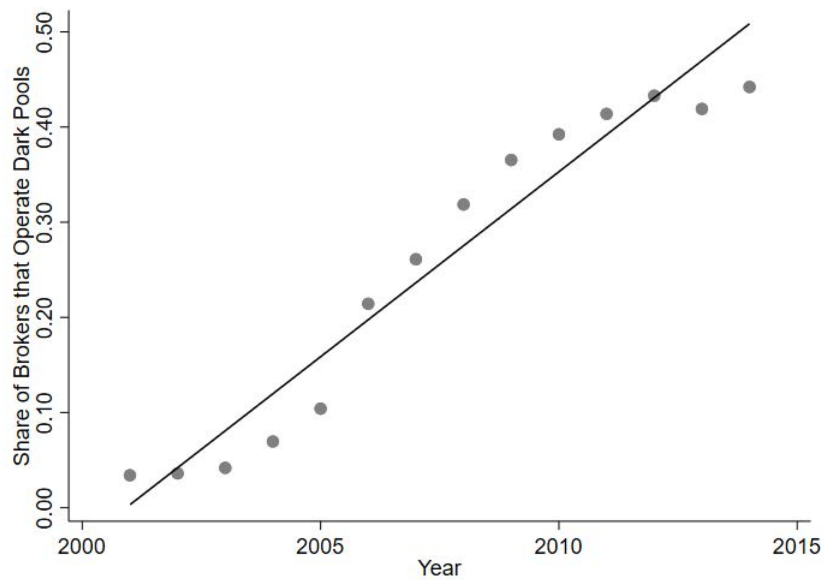

Note: Figure A1 displays a binned scatter plot of the share of brokers that operate alternative trading systems/dark pools over time. Observations are at the investor-by-broker-by-sector-by-month level which is the unit of observation in our main analysis.

Figure A2: Distribution of Investor-Broker-Time Fixed Effects

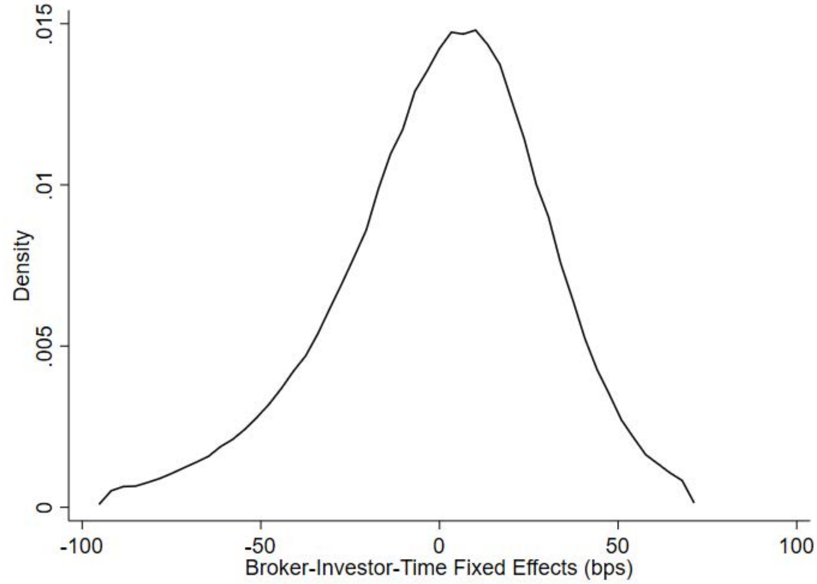

Note: Figure A2 displays the distribution of estimated investor-broker-time fixed effects corresponding to column (5) in Table 2. The fixed effects are scaled by $1/\alpha$, which puts the fixed effects in terms of basis points (bps) rather than utils. Observations are averaged at the investor-by-broker-by-sector-by-month level which is the unit of observation in our main analysis.

Table A1: Broker Choice: Alternative Market Definition (Investor-by-Year)

	(1)	(2)	(3)
Fees	-546***	-768***	-688***
	(9.99)	(16.3)	(14.3)
Price Impact	7.40***	-6.89	14.8
	(1.55)	(29.4)	(41.1)
Research:			
Number of Analysts	0.17***	0.13***	0.036***
	(0.0091)	(0.0092)	(0.010)
Number of Top Rated Analysts	0.40***	0.57***	0.20***
	(0.070)	(0.061)	(0.062)
Information:			
Eigenvector Centrality	4.24***	3.30***	0.70***
	(0.43)	(0.40)	(0.24)
Informed Broker	2.28***	1.83***	0.55***
	(0.073)	(0.072)	(0.061)
Investor×Time Fixed Effects	X	X	X
Broker Fixed Effects			X
IV (Commissons & Price Impact)		X	X
Cragg Donald F-Statistic for IV		395	116
Observations	106,049	84,887	84,870
R-squared	0.379	0.340	0.405

Note: Table A1 displays the estimation results corresponding to our discrete choice broker model (eq. 8). Here we define the market at the investor-by-year level such that the unit of observation is at the investment manager-by-broker-year (6-digit GICS) level over the period 2001-2014. Each independent variable is described in detail in Section 4.2. We measure fees in percentage terms relative to the value of the transaction. As described in the text, we instrument for fees using the average historical fee charged by the broker in terms of cents per share divided by the share price of the stock being traded. We instrument for price impact using the lagged price impact to account for measurement error. Standard errors are clustered at the broker-by-year level and are reported in parentheses. The symbols *,**, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table A2: Broker Choice and Alternative Trading Systems

	(1)	(2)	(3)
Fees	-151***	-412***	-402***
	(4.47)	(7.96)	(7.14)
Alternative Trading Systems (ATS)	0.10***	0.085***	0.12***
	(0.030)	(0.031)	(0.028)
Sector×Investor×Time Fixed Effects	X	X	X
Broker Fixed Effects			X
IV (Commissons & Price Impact)		X	X
Cragg Donald F-Statistic for IV		6,188	2,878
Observations	6,187,028	5,607,856	5,607,852
R-squared	0.302	0.270	0.299

Note: Table A2 displays the estimation results corresponding to our discrete choice broker model (eq. 8). The unit of observation is at the investment manager-by-broker-by-month-by-sector (6-digit GICS) level over the period 2001-2014. Alternative Trading System (ATS) is a dummy variable indicating that the broker operates an ATS in the corresponding year. We measure fees in percentage terms relative to the value of the transaction. As described in the text, we instrument for fees using the average historical fees charged by the broker in terms of cents per share divided by the share price of the stock being traded. Other controls include: Price Impact, Number of Research Analysts, Number of Top Research Analysts, Broker Eigenvector Centrality, and Informed. We instrument for price impact using the lagged price impact to account for measurement error. Standard errors are clustered at the broker-by-year level and are reported in parentheses. The symbols *,**, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table A3: Autocorrelation of Broker Characteristics (12-month)

	(1)	(2)	(3)	(4)	(5)	(6)
Fees	0.42***					
	(0.0054)					
Price Impact:		0.033***				
		(0.0025)				
Research:						
Number of Analysts			0.88***			
			(0.015)			
Number of Top Rated Analysts				0.79***		
				(0.017)		
Information:						
Eigenvector Centrality					0.35***	
					(0.020)	
Informed Broker						0.39***
						(0.010)

Note: Table A3 displays the 12-month autocorrelation of the broker controls in our baseline analysis.

Table A4: Broker Choice - Instrumenting for Research Coverage

	(1)	(2)	(3)	(4)
Fees	-532***	-493***	-492***	-420***
	(11.0)	(9.95)	(9.87)	(7.59)
Price Impact:	-84.8***	-25.1	-18.1	-20.0
	(31.0)	(24.7)	(19.7)	(19.7)
Research:				
Number of Analysts	0.082***	0.045***	0.052***	0.047***
	(0.0059)	(0.0055)	(0.0036)	(0.0025)
Number of Top Rated Analysts	0.19***	0.067***	0.059***	0.079***
	(0.023)	(0.012)	(0.0080)	(0.0069)
Information:				
Eigenvector Centrality	1.21***	0.40***	0.13**	0.59***
	(0.085)	(0.055)	(0.056)	(0.037)
Informed Broker	0.29***	0.099***	0.079***	0.071***
	(0.017)	(0.0081)	(0.0049)	(0.0044)
Sector×Investor×Time Fixed Effects	X	X	X	X
Broker Fixed Effects		X		
Broker×Time Fixed Effects			X	
Broker×Investor×Time Fixed Effects				X
IV (Commissons, Price Impact, & Research)	X	X	X	X
Cragg Donald F-Statistic for IV	1,505	631	285	240
Observations	2,382,835	2,382,810	2,380,391	2,320,361
R-squared	0.217	0.276	0.308	0.583
Value of Research:				
Value of an Additional Analyst (bp)	1.54	0.91	1.06	1.12
Value of an Additional Top Analyst (bp)	5.11	2.27	2.26	3.00

Note: Table A4 displays the estimation results corresponding to our discrete choice broker model. The unit of observation is at the investment manager-by-broker-by-month-by-sector (6-digit GICS) level over the period 2001-2014. Each independent variable is described in detail in Section 4. We measure fees in percentage terms relative to the value of the transaction. As described in the text, we instrument for fees using the average historical fee charged by the broker in terms of cents per share divided by the share price of the stock being traded. We instrument for price impact using the lagged price impact to account for measurement error. We also instrument Number of Analysts and Number of Top Rated Analysts using their lagged values (12-month lags). Standard errors are clustered at the broker-by-year level and are reported in parentheses. The symbols *,**, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.
In the bottom half of the table, we interpret the coefficient estimate. We compute the value of each independent variable as the ratio of the coefficient of interest divided by an investor's sensitivity with respect to fees ($-\alpha$). For example, we calculate the value of an analyst in column (1) as 10,000 $\times$ 0.082/532 = 1.54bps.

Table A5: Broker Choice and Large Implicit Trading Costs (Large Price Impact)

	(1)	(2)	(3)	(4)	(5)
Fees	-152***	-400***	-401***	-402***	-373***
	(4.46)	(8.07)	(7.17)	(7.12)	(5.49)
Large Price Impact	0.031***	-1.35***	-0.61***	-0.63***	-0.62***
	(0.0066)	(0.21)	(0.15)	(0.17)	(0.17)
Sector×Investor×Time Fixed Effects	X	X	X	X	X
Broker Fixed Effects			X		
Broker×Time Fixed Effects				X	
Broker×Investor×Time Fixed Effects					X
IV (Commissons & Price Impact)		X	X	X	X
Cragg Donald F-Statistic for IV		4,844	1,899	720	664
Observations	6,187,028	5,607,856	5,607,852	5,607,323	5,487,670
R-squared	0.302	0.187	0.281	0.298	0.572
Mean Elasticity with Respect to Fees	0.18	0.47	0.47	0.47	0.44
Value of a Large Price Impact (bp)	2.04	-33.75	-15.21	-15.63	-16.62

Note: Table A5 displays the estimation results corresponding to our discrete choice broker model (eq. 8). The unit of observation is at the investment manager-by-broker-by-month-by-sector (6-digit GICS) level over the period 2001-2014. Each independent variable is described in detail in Section 4.2. We measure fees in percentage terms relative to the value of the transaction. As described in the text, we instrument for fees using the average historical fee charged by the broker in terms of cents per share divided by the share price of the stock being traded. Large Price Impact measures whether the price impact was greater than 0.25%. The median and average price impact of a Large Price Impact trade is 0.46% and 0.65%. We instrument for Large Price Impact using an indicator variable that indicates whether the lagged Price Impact (12 month rolling average) of the broker in that sector is greater than 0.25%. Standard errors are clustered at the broker-by-year level and are reported in parentheses. The symbols *,**, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

In the bottom half of the table we interpret the coefficient estimates. Elasticity of demand is calculated as the average of $-\alpha * (1-s) * fee$. We compute the value of each independent variable as the ratio of the coefficient of interest divided by an investor's sensitivity with respect to fees ($-\alpha$). For example, we calculate the value of avoiding a large price impact in column (2) -1.35/400 = -34bps.