Analyst Labor Market Allocation and Price Efficiency: A Structural Approach*

Mingze Ma† Hongwei Mo‡

June 2026

Abstract

We study how the labor-market allocation of financial analysts affects stock price efficiency. Using brokerage mergers and senior analyst retirements as quasi-natural experiments, we document causal effects of brokerage platform matching and team composition on analyst forecast informativeness. However, analyst job mobility remains low and is declining over time, which may undermine the quality of matching. To quantify the impact of the analyst labor market on price efficiency, we develop and estimate a structural model that unifies forward-looking broker choice subject to idiosyncratic preferences and mobility costs, Nash-in-Nash bargaining under oligopolistic competition, and heterogeneous analyst-broker matching. Our results suggest that labor-market frictions limit price efficiency primarily by impeding analyst mobility toward mid- and lower-tier brokers that rely more heavily on high-quality analyst teams. Reducing job-switching costs to one tenth of their estimated level, for example, raises analysts' aggregate price impact by 8.25%.

***JEL classification**: C78, G14, G24, J44, J62, L13*

***Keywords**: Sell-Side Analyst, Labor Market Frictions, Search and Matching, Price Efficiency, Structural Estimation*

*We are grateful to Francois Bares, Hui Chen, Qingyu Chen, Niels Johannesen, Alok Kumar, Seung Joo Lee, Shumiao Ouyang, René Stulz, and Emmanuel Yimfor for helpful comments and suggestions. We thank Kaiyue Qian for excellent research assistance. All errors are our own.

†Saïd Business School, University of Oxford. Email: mingze.ma.dphil@said.oxford.edu.

‡Saïd Business School, University of Oxford. Email: hongwei.mo@sbs.ox.ac.uk.

1 Introduction

Financial markets rely on specialized intermediaries for the efficient flow of information. Sell-side equity analysts are among the central information producers in capital markets. They collect firm-specific data, interpret information, and communicate their views to investors through research reports, earnings forecasts, and stock recommendations (Bradshaw et al., 2017). A large literature establishes that analyst research moves prices and shapes market efficiency.[1] Yet analysts do not produce information in a vacuum. They are embedded in brokerage platforms that differ enormously in team composition, research infrastructure, and client networks. As a result, the same analyst may produce very different research depending on where she works. This makes the labor market that allocates analysts across brokers part of the information-production technology of capital markets. However, we know much less about how this market works or whether its allocation of talent affects how efficiently information enters prices.

This gap matters. In a frictionless labor market, analyst talent would flow to the brokers and teams where it generates the greatest informational value. In practice, however, various labor market frictions such as imperfect competition and mobility constraints may distort this allocation.[2] This paper asks how these frictions shape the labor market allocation of financial analysts, and how much this allocation affects stock price efficiency. We answer these questions by first providing reduced-form evidence that analyst allocation across brokers matters, and then developing a structural model of the analyst labor market to quantitatively link this allocation to analyst research quality and stock price efficiency.

We begin by establishing the economics of analyst-broker and analyst-team matching. First, *platform quality and analyst ability are substitutes.* That is, a stronger platform benefits low-ability analysts more than high-ability ones. We exploit brokerage mergers and acquisitions (M&As) that move analysts from target to larger acquiring brokers. These events are driven by business-level strategy and are therefore plausibly exoge-

[1] See, among others, Womack (1996), Hong et al. (2000b), Barber et al. (2001), Gleason and Lee (2003), Loh and Stulz (2011), Kelly and Ljungqvist (2012), Huang et al. (2014), and Loh and Stulz (2018). For a survey of analyst forecasts and their asset-pricing implications, see Kothari et al. (2016).

[2] Industry evidence points to persistent matching frictions in financial-services labor markets. In the 2025 Robert Half Salary Guide, 93 percent of hiring managers in financial services report difficulty finding qualified candidates (Robert Half, 2025).

nous to analysts. Using a stacked difference-in-differences (DiD) design, we compare post-merger changes in the informativeness of forecast revisions for low- versus high-experience movers around the same event, within analyst-firm pairs. After moving to a larger platform, junior analysts experience a significantly larger increase in the market reaction to their revisions than senior analysts do: A one-standard-deviation increase in the revision raises the announcement-window cumulative abnormal return (CAR) by 0.78 percentage points more for juniors.

Second, *analysts are complements within teams.* We identify retirements of highly experienced analysts from LinkedIn data and use them as plausibly exogenous shocks to team composition, comparing remaining junior analysts who cover the departing analyst's industry (treated) with other junior colleagues at the same brokerage who do not (control). Following a senior departure, the informativeness of forecast revisions falls sharply for treated analysts relative to controls: A one-standard-deviation revision generates a 1 percentage point smaller announcement-window CAR. This result is consistent with complementarity within analyst teams, whereby senior analysts generate positive informational spillovers for their junior colleagues.

The relevance of labor-market matching is sharpened by a striking secular trend: *Analyst job mobility has fallen markedly*, from 14.56% in 2001 to about 4.06% by 2024, and *declines steeply with analyst experience*. A labor market that is becoming less mobile, especially among its most skilled participants, may allocate information-producing talent less effectively. Whether this matters for price efficiency, and through which margins, is a quantitative question that reduced-form evidence alone cannot answer.

To characterize the allocation of financial analysts across brokers and quantify the effects of different labor-market frictions on analysts' aggregate information production, we develop a structural model that incorporates multiple labor-market mechanisms and sources of heterogeneity. Figure 1 provides an overview of the model. In the model, brokers provide profitable brokerage services to clients under oligopolistic competition. To attract greater client demand and extract more surplus, brokers compete in the labor market to form analyst teams and combine them with their own platform resources.

To capture heterogeneity in the fit between analysts and brokers, we allow for flexible matching mechanisms between brokers and analyst teams, as well as among different analyst types within teams. Once matched, the quality of analysts' information production

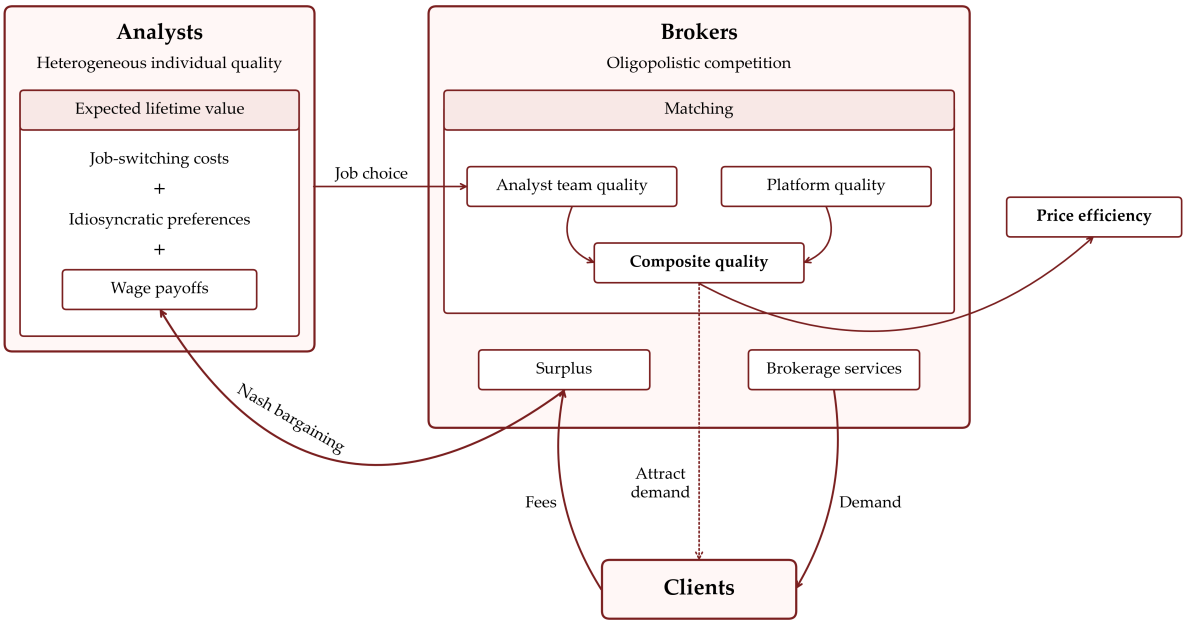


Figure 1. Model Overview.

depends jointly on broker resources and team ability. We model this composite quality as the scale parameter in clients' Fréchet preference distribution, thereby linking analyst allocation to broker surplus.

Given equilibrium wage payoffs, analysts make forward-looking broker choices in each period to maximize expected lifetime value, taking into account wage flows, future job-switching opportunities, mobility costs, and Gumbel-distributed idiosyncratic preference shocks. The properties of the Gumbel distribution allow us to characterize the law of motion for analyst allocation across brokers.

Equilibrium wage payoffs, in turn, are determined by bargaining between brokers and analysts. Under standard assumptions, we characterize this process as Nash bargaining embedded in the Nash equilibrium of a non-cooperative game. That is, each broker-analyst pair sets wages to maximize the exponentiated product of the two parties' surpluses from agreement relative to disagreement, taking the wages set by other brokers as given. Because brokers' surplus from analysts depends on profits from brokerage services and the quality of analyst-broker matching, while analysts' surplus from brokers depends on the expected value of alternative job choices, all components of the model are endogenously interconnected.

For quantification, we focus on a steady-state equilibrium with a stationary allocation of analysts across brokers and provide an efficient damped fixed-point iteration algorithm to solve for it. We then combine Poisson pseudo-maximum likelihood (PPML) and simulated method of moments (SMM) estimation to recover the model parameters that fit the analyst allocation data.

We validate our estimation results in three ways. First, the model-implied analyst allocation matches the targeted moments well. Second, the model-implied total surplus attributable to analysts is significantly positively correlated with manually collected brokerage service revenue, which is not targeted in the SMM estimation. Finally, and most importantly, the model-implied broker composite quality is strongly positively associated with the data-based measure of broker-level aggregate informativeness of analyst reports, which validates the model's ability to capture the impact of analyst allocation on price efficiency.

The parameter estimates confirm the two matching patterns from the reduced-form evidence. Platform quality and analyst-team quality are substitutes. A strong broker with established client networks, superior information access, and reputational capital gains little marginal benefit from adding high-quality analysts. Brokers without such resource advantages, by contrast, must rely on high-quality analyst teams to compete. At the same time, analysts of different types are complementary within teams. A more balanced team generates stronger synergies, making its value greater than the sum of its individual members. We also estimate a low elasticity of client demand with respect to brokerage-service prices, which reinforces the importance of high-quality sell-side research, rather than low commission fees, in attracting clients and extracting surplus. In the model, however, a few brokers with exceptionally strong platform quality attract disproportionately many clients with little reliance on high-quality analyst teams, resulting in a highly concentrated market structure.

We then use the estimated model to quantify how each labor-market mechanism affects analysts' aggregate contribution to price efficiency. Taken together, the estimation results deliver the model's central insight: Labor-market frictions limit price efficiency mainly by impeding the flow of analysts toward the mid- and lower-tier brokers that rely more heavily on high-quality analyst teams. Any shock that facilitates this flow thus improves the efficiency of analyst allocation, and vice versa. For example, reducing job-

switching costs to one tenth of their estimated level raises analysts' aggregate contribution to price efficiency by 8.25%, and equalizing analysts' bargaining power with brokers' increases the aggregate impact by 3.46%. Intensifying price competition in the market for brokerage services, by contrast, reduces analysts' aggregate informativeness, as brokers shift their focus away from analyst talent toward fee competition.

Finally, the matching technology itself is first-order. Removing the substitution between platform quality and team quality raises aggregate informativeness by 22.51%, the largest effect among our counterfactuals. Dominant platforms produce high-quality research largely through their own informational resources, leaving them with limited incentives to compete intensively for analyst-team quality in the labor market. When the substitution is removed, strong teams at top brokers are no longer discounted by platform strength, and broker success depends more directly on analyst-team quality. In contrast, removing within-team complementarity lowers aggregate informativeness by 4.64%, a much smaller effect, because team synergies matter primarily for mid- and lower-tier brokers.

These quantitative results carry direct policy implications for sell-side research. Because the substitution between platform quality and analyst-team quality is the largest drag on efficiency, a central policy priority is to improve transparency in research pricing. Requiring clients to evaluate analyst research separately from the platform through which it is delivered would strengthen the link between payments and analyst-team quality, giving brokers stronger incentives to compete for analyst talent. This speaks to the ongoing debate over the bundling of research and execution payments across the US, EU, and UK (FCA, 2024, 2025; EU, 2024).

Three further policy margins also follow from our counterfactuals. First, better calibrated non-compete agreements would reduce frictions to analyst mobility (Starr et al., 2021), allowing talent to reach the brokers where its marginal contribution is greatest. Second, lowering the fixed costs borne by small research providers and expanding their retail access would strengthen competition in research provision (Kent, 2023; Drake et al., 2023), especially where the gains from analyst mobility are largest. Third, greater wage transparency and portable records of analyst performance would make analyst value more visible to rival employers, strengthening the bargaining position of experienced analysts.

Related literature Our work builds on a large literature showing that sell-side research contributes to price discovery and market efficiency. In particular, analyst output has value and moves prices (e.g., Womack, 1996; Barber et al., 2001; Gleason and Lee, 2003), analyst coverage reduces information asymmetry (Kelly and Ljungqvist, 2012), and limited coverage slows the incorporation of bad news into prices (Hong et al., 2000b). A related literature examines the characteristics that make analysts and their research more informative, including industry experience (Bradley et al., 2017), all-star status (Stickel, 1992), and the qualitative content of research reports (Huang et al., 2014). This literature, however, focuses on the information production of individual analysts in reduced-form frameworks while taking the allocation of analysts across brokers as given. We contribute to this literature by shifting the object of interest to labor-market mechanisms, the resulting endogenous equilibrium allocation of analysts, and the aggregate consequences of this allocation for stock prices. To our knowledge, this is the first paper to model the labor market for sell-side analysts and link the allocation of information producers to aggregate price efficiency. We show that labor-market allocation is a determinant of analyst informativeness, characterize the underlying mechanisms in equilibrium, and quantify the aggregate efficiency loss from misallocation.

Another strand of literature documents that an analyst's productivity depends not only on individual ability but also on her employer and colleagues. Brokerage resources influence forecast accuracy (Clement, 1999), career concerns and brokerage hierarchies affect analysts' research incentives and outputs (Hong et al., 2000a; Hong and Kubik, 2003), and peer effects among analysts contribute to information production (Huang et al., 2022; Phua et al., 2023; Mo, 2024). This literature shows that labor-market matching matters but leaves the exact matching mechanisms largely unspecified. We use both reduced-form and structural analyses to characterize these mechanisms: Broker platform quality and analyst ability are substitutes, while analysts within a team are complements. We also quantify how job mobility affects the quality of analyst-broker matching, an issue that has become increasingly important given the sharp decline in analyst mobility over time and the broader fall in labor-market fluidity (Molloy et al., 2016).

We also connect to a broader literature linking labor-market dynamics to financial-market outcomes, including assignment models of how the matching of talented agents to firms shapes valuations and price efficiency (Chang and Hong, 2017, 2019) and evi-

dence that labor-market search frictions are priced in the cross-section of equity returns (Kuehn et al., 2017). We share the spirit of this literature that the allocation of labor matters for financial outcomes, but study a different market and mechanism. Rather than focusing on the matching of issuers and underwriters in primary markets, we study the labor market for sell-side analysts in secondary markets. Rather than treating labor frictions as a source of priced risk in stock returns, we trace how these frictions distort the production of public information.

To the best of our knowledge, our structural model is the first to unify workers' dynamic forward-looking employer choice subject to idiosyncratic preference shocks and mobility costs, "Nash-in-Nash" wage bargaining alongside oligopolistic labor and product markets, and heterogeneous worker-firm matching into a tractable and quantifiable framework.

Our characterization of analysts' job choices follows the dynamic discrete choice models of Artuç et al. (2010), Caliendo et al. (2019), and Kleinman et al. (2023). In those models, however, labor-market imperfections arise only from workers' mobility costs across sectors, so wages are general-equilibrium objects that apply to all firms within a given sector. We contribute to this framework by incorporating oligopolistic labor-market competition and worker-firm wage bargaining. By making wage determination granular at the firm level, our model generates rich wage heterogeneity across analyst-broker pairs, which helps match the observed allocation of analyst labor and more accurately quantify counterfactual effects.

The "Nash equilibrium in Nash bargains" is widely used in applied industrial organization to characterize surplus division in oligopolistic markets, including healthcare (Gowrisankaran et al., 2015), telecommunications (Crawford et al., 2018), and apparel (Eaton et al., 2026). Collard-Wexler et al. (2019) provide microfoundations for these applications. We extend this literature not only by applying this theoretical structure to the labor market for financial analysts but also by linking labor-market bargaining to surplus generated in an oligopolistic product market: profitable brokerage services. Our model is also applicable to other labor markets in which workers produce intangible outputs that attract demand in profitable product markets, such as the labor markets for R&D and academic activities.

For labor-market matching, we use a nested constant elasticity of substitution (CES)

structure to jointly characterize matching between brokers and analyst teams and matching among analysts within teams. This differs from the existing literature, which typically studies either worker-firm matching (e.g., Eeckhout and Kircher, 2011; Menzio and Shi, 2011; Lise and Robin, 2017; Bonhomme et al., 2019; Lamadon et al., 2022) or within-team matching (Herkenhoff et al., 2024). We remain agnostic about the parameters governing the matching technology and estimate them to fit analyst allocation data. Our estimates imply substitutability between brokers and analyst teams, which differs from the worker-firm complementarities widely documented using economy-wide labor-market data (e.g., Lise and Robin, 2017; Bonhomme et al., 2019; Lamadon et al., 2022). We argue that this matching mechanism may be specific to financial services, where platform resources and team ability are highly substitutable in producing intangible outputs that attract clients. At the same time, our finding of complementarity among analysts across experience groups echoes the hierarchical organization of knowledge production proposed by Garicano (2000) and the spatial complementarity between high- and low-skill workers documented by Eeckhout et al. (2014).

Finally, although our focus is on how labor market allocation affects analysts' information production and aggregate price efficiency, our model also captures several patterns consistent with the labor economics literature, including occupation-specific human capital accumulated through experience (e.g., Kambourov and Manovskii, 2009; Bagger et al., 2014), worker heterogeneity in bargaining power (e.g., Cahuc et al., 2006; Caldwell et al., 2026), and wage dispersion driven by worker effects and firm effects, especially the former (e.g., Abowd et al., 1999; Postel-Vinay and Robin, 2002; Bonhomme et al., 2019).

Layout The remainder of the paper is organized as follows. Section 2 presents the institutional background and motivating evidence. Section 3 develops the structural model. Section 4 reports the estimation, validation, and counterfactuals. Section 5 concludes.

2 Motivating Evidence

This section presents institutional background and empirical evidence that motivate our structural analysis. The central premise is that the informativeness of analyst research depends not only on individual analyst ability, but also on the brokerage platform and

analyst team in which the analyst is embedded. We first describe the labor market for sell-side analysts. We then exploit brokerage M&As that reassign analysts to larger platforms to examine how analyst-broker matching affects the informativeness of analyst research. Next, we use the departure of highly experienced analysts as shocks to team composition to study interactions among analysts within teams. Finally, we document that cross-broker job mobility is low and has declined substantially over time, particularly among highly experienced analysts.

2.1 Data

We obtain analysts' annual earnings forecasts from the Institutional Brokers' Estimate System (IBES). We use the unadjusted detail file and adjust forecasted EPS using the split-adjustment factor from the Center for Research in Security Prices (CRSP).[3] Firm characteristics and stock returns are from Compustat and CRSP. We retain only common stocks (SHRCD = 10 or 11) listed on the NYSE, AMEX, and Nasdaq (EXCHCD = 1, 2, or 3). The sample period runs from January 1995 to December 2024.[4] Following prior literature (Clement, 1999; Clement et al., 2007; Harford et al., 2019), we restrict the sample to earnings forecasts issued during the first 11 months of a fiscal year, ensuring a minimum forecast horizon of 30 days. For analyst employment histories, we manually collect supplementary data using LinkedIn and BrokerCheck.[5] The sample contains 16,551 distinct analysts from IBES, of which we are able to match 5,242 on LinkedIn.

2.2 Institutional Background: The Labor Market for Sell-Side Analysts

Sell-side analysts are professional information intermediaries employed by brokerage houses and investment banks. Their work involves less observable activities, such as collecting and interpreting firm-specific information using industry expertise, as well as

[3] The IBES Adjusted Detail file provides EPS forecasts that are adjusted for stock splits and rounded to two decimal places. Diether et al. (2002) and Payne and Thomas (2003) note that this rounding can introduce adjustment bias. Detailed guidance on adjusting EPS forecasts is available from Wharton Research Data Services (WRDS).

[4] We begin the sample in 1995 because the quality of analyst data improves after that year (Clement et al., 2011; Hilary and Hsu, 2013).

[5] We first manually collect analyst information from BrokerCheck and cross-validate it using LinkedIn. BrokerCheck is provided by the Financial Industry Regulatory Authority (FINRA) and is available at https://brokercheck.finra.org/.

communicating with institutional clients. These less visible dimensions of work also shape analyst career outcomes (Brown et al., 2015; Bradley et al., 2017). Through these information-intermediation activities, analysts create indirect value for brokers: Their research attracts buy-side attention, generates trading commissions, and supports broader brokerage and investment-banking business (Brown et al., 2015).

Analysts' information production depends critically on their brokerage platforms. Much of the economic value of analyst research is realized only when information is distributed through a brokerage platform and used by institutional clients or other market participants. Accordingly, the informativeness of analyst research depends not only on individual ability but also on employer resources, including access to management (Green et al., 2014), client relationships (Brown et al., 2015), and the prestige of brokerage houses (Hong and Kubik, 2003). Brokers, in turn, rely on talented and reputable analysts to attract institutional clients and support trading and research-related business (Jackson, 2005; Gokkaya et al., 2023). The analyst labor market can therefore be viewed as a two-sided matching environment in which heterogeneous analysts are matched with heterogeneous brokerage platforms. The quality of this match matters because it can shape the informativeness of analyst research and the efficiency with which information is incorporated into prices.

2.3 Brokerage Platform Matters

In this section, we examine whether and how brokerage platforms affect the informativeness of analyst research. The mechanism underlying analyst-broker matching is ambiguous. If analyst ability and platform quality are substitutes, lower-quality analysts may benefit more from high-quality broker resources. If they are complements, high-quality broker resources may be exploited more effectively by high-ability analysts in information production. If, instead, analysts and brokers are independent, brokerage platforms should be unrelated to the informativeness of analyst research.

The key challenge in the empirical design is that analyst-broker matches observed in the data are equilibrium outcomes distorted by various labor-market frictions. To address this endogeneity concern and identify the effect attributable to brokerage platforms, we exploit brokerage M&As that move analysts from target brokers to acquiring brokers

(Hong and Kacperczyk, 2010; Kelly and Ljungqvist, 2012). These events are plausibly exogenous with respect to analyst-broker matching because they are primarily driven by business-level strategic considerations (Derrien and Kecskés, 2013). We focus on events in which the acquiring broker is larger than the target broker.[6] The final sample contains 20 M&A events from 1997 to 2020, covering 188 moving analysts and 1,973 firms. Details of the event construction and mover definition are provided in Table A1.

We examine the analyst-broker matching mechanism by splitting analysts into two groups based on their work experience, a proxy for analyst quality (Clement, 1999). An analyst is classified as low-experience (treatment group) if her pre-event experience is below the annual median.[7] The identifying assumption is that, absent the M&A event, low- and high-quality moving analysts would not have experienced differential changes in the price informativeness of their forecast revisions around the same M&A event. Specifically, we estimate the following stacked DiD regressions (Baker et al., 2022):

$$\begin{aligned} CAR_{j,f,\tau,e} = {} & \beta_1 \times SignedRevision_{j,f,\tau} + \beta_2 \times Post_{e,\tau} + \beta_3 \times Treatment_{j,e} \times Post_{e,\tau} \\ & + \beta_4 \times SignedRevision_{j,f,\tau} \times Post_{e,\tau} \\ & + \beta_5 \times SignedRevision_{j,f,\tau} \times Treatment_{j,e} \\ & + \beta_6 \times SignedRevision_{j,f,\tau} \times Treatment_{j,e} \times Post_{e,\tau} \\ & + \alpha_{j,f,e} + \delta_t + \varepsilon_{j,f,\tau,e}. \end{aligned} \tag{1}$$

The dependent variable is the three-day CAR, adjusted by the CRSP value-weighted index, around analyst j's forecast announcement for firm f on date τ in merger event e. $SignedRevision_{j,f,\tau}$ is the signed forecast revision, scaled by the stock price two days before the announcement. Its coefficient captures the sensitivity of stock prices to analyst forecast revisions and is therefore interpreted as a measure of the price informativeness of analyst research (Gleason and Lee, 2003; Hilary and Hsu, 2013).[8] $Post_{e,\tau}$ equals one

[6] Broker size is measured by the number of unique analysts who issued at least one forecast during the 12 months preceding the event date.

[7] Analyst experience is defined as the number of years from the first forecast documented in IBES to the current forecast.

[8] This measure follows a large literature that evaluates the information content of analyst research from the contemporaneous stock-price response to forecast revisions and other outputs. Forecast revisions convey news about expected cash flows, and revisions that elicit stronger directional price responses are interpreted as containing more value-relevant information (Lys and Sohn, 1990; Gleason and Lee, 2003; Asquith et al., 2005; Bradley et al., 2014)

after the analyst moves to the acquiring broker, and $Treatment_{j,e}$ identifies low-experience moving analysts. We include analyst-event-firm fixed effects ($\alpha_{j,f,e}$) and year-month fixed effects (δ_t). These fixed effects absorb all cross-sectional heterogeneity in analyst ability, firm characteristics, deal-specific factors, analyst-firm optimism (Kumar et al., 2022), as well as aggregate time-series variation. Identification thus comes from comparing the change in forecast informativeness after the move for low-experience analysts relative to high-experience analysts, while holding fixed the same analyst-firm pair around the same M&A event. Standard errors are clustered by M&A events (Abadie et al., 2023).

Figure 2 reports the results. In Panel A, we run separate regressions for six different time windows around the event, ranging from 6 to 21 months on each side. In the 12-month window, for example, the coefficient on *SignedRevision* $\times$ *Treatment* $\times$ *Post* is 0.402 and statistically significant at the 1% level. This indicates that, after moving to a larger acquiring broker, junior analysts experience a larger increase in the market reaction to their forecast revisions than senior analysts. For the economic magnitude, a one-standard-deviation increase in *SignedRevision* is associated with an additional 0.78 percentage-point increase in announcement-window CAR for low-experience analysts after the move (0.402×0.0194). The effect is also positive and significant using 15-, 18-, and 21-month windows (0.458, 0.517, and 0.530), but is smaller and statistically insignificant using 6- and 9-month windows. This pattern suggests that the effect is concentrated in horizons long enough for the platform shock to affect analyst informativeness.

A natural concern with this design is that low- and high-experience movers might have been on diverging informativeness trajectories even in the absence of the M&A. This could arise, for instance, if less experienced analysts improve more rapidly at an earlier stage of the learning curve. To address this concern, Panel B of Figure 2 reports a dynamic event-study specification for the 12-month window, plotting treatment effects at three-month intervals relative to the move, with the period immediately preceding the move (month -3) as the reference point. The estimates at months -12, -9, and -6 are economically small, statistically indistinguishable from zero, and show no drift toward the post-event effect. The differential treatment effect emerges only after the move, growing from 0.150 at month +3 to 1.051 at month +12. This timing is difficult to reconcile with the learning-curve story, which would predict a smooth trajectory spanning the move rather than a discrete break at the event. As an additional falsification test, we assign each event a

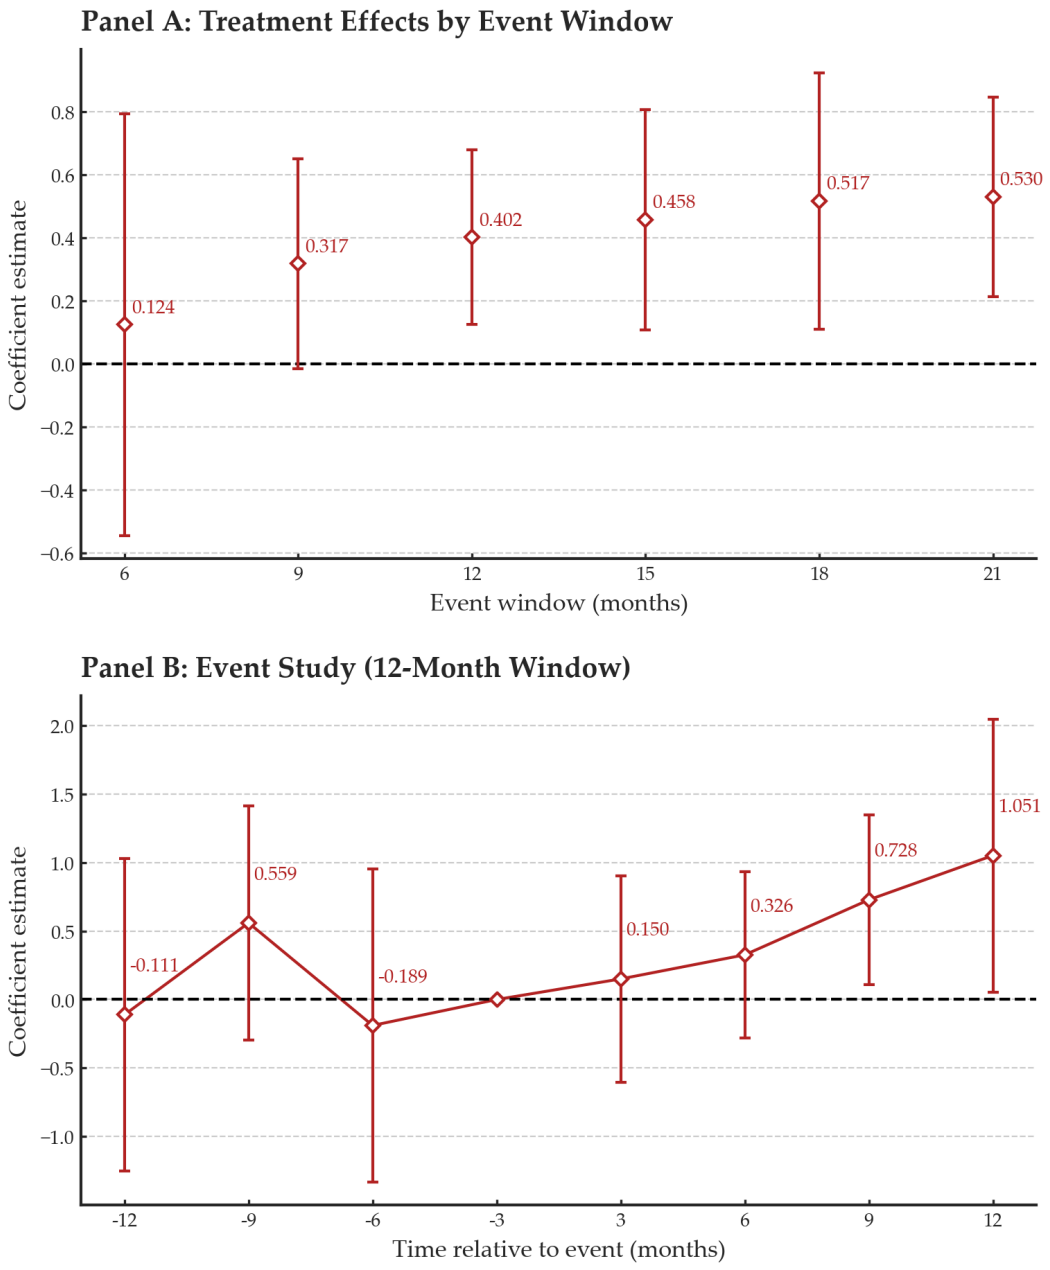


Figure 2. Brokerage M&As and Analyst Forecast Informativeness.

Notes: This figure reports stacked DiD estimates of the effect of brokerage M&As on the price informativeness of analyst forecast revisions, comparing low-experience movers (treatment) with high-experience movers (control). Informativeness is measured by the sensitivity of the three-day announcement CAR to the signed forecast revision, and the reported coefficient is on the triple interaction term *SignedRevision* × *Treatment* × *Post* in equation (1). An analyst is assigned to the treatment group if her pre-event tenure is below the annual median. We exclude forecasts made within the one-month window on either side of the event date. All specifications include analyst-event-firm and year-month fixed effects, with standard errors clustered by M&A event. Panel A plots the coefficient estimates from six separate regressions, each using a symmetric pre/post estimation window of the length shown on the horizontal axis. Panel B reports a dynamic event-study specification estimated on the 12-month window, plotting the triple-interaction coefficient by quarter relative to the move, with the quarter immediately preceding the move as the reference point. Point estimates are labeled, and vertical bars denote 95% confidence intervals.

pseudo-event date 24 months before the actual event and re-estimate the specification using analysts who remain at the target brokerage. Figure A1 shows that the placebo estimates are statistically indistinguishable from zero across event windows and exhibit no systematic post-event pattern.

Overall, the evidence shows that larger brokerage platforms raise the informativeness of low-experience analysts more than that of their high-experience peers. This pattern suggests that broker resources and individual analyst ability are substitutes rather than complements in the production of informative research, since the marginal value of a stronger platform is greatest for analysts who have the least to draw on themselves. More broadly, it indicates that brokerage platforms, not just individual ability, shape the information that analysts' forecasts convey to the market.

2.4 Analyst Team Composition Matters

We next examine how team composition affects the informativeness of analyst research. Analysts do not work in isolation.[9] Within a brokerage house, analysts may benefit from colleagues' industry knowledge, client interactions, management access, and coverage of economically related firms (Huang et al., 2022; Mo, 2024). These interactions imply that the value of one analyst depends on the composition of the broader analyst team. Experienced analysts, in particular, may supply the industry expertise, institutional knowledge, and client credibility that help their colleagues produce more informative research.

We exploit retirements of senior analysts as plausibly exogenous shocks to team composition. Using analysts' LinkedIn profiles, we infer each analyst's age and examine their employment history, from which we compute years of financial-industry experience. We define a senior analyst retirement as the permanent disappearance from IBES of an analyst who, at the time of departure, is older than 55 and has at least 20 years of financial-industry experience. We further require that the analyst is not subsequently observed moving to another employer or industry, which ensures that the event captures exit from the labor force rather than a job change. Under these stringent requirements, we are able to identify 45 retirement events spanning 2002 to 2023.

[9] Fang and Hope (2021) study hand-collected research reports and show that team-based production is pervasive in sell-side research. They document that analyst teams issue a large share of earnings forecasts and that team-issued forecasts are more accurate than forecasts issued by individual analysts.

We study the impact of these departures on the forecast informativeness of remaining junior analysts at the same brokerage house. We restrict the remaining analysts to those with fewer than five years of experience.[10] Treated analyst-firm observations are those in which the junior analyst covers a firm in the same Fama-French 30 industry as the departing senior analyst. Control observations are remaining junior analyst-firm observations at the same brokerage house that are not exposed to the departing analyst's industry.[11]

This design helps isolate team-composition effects from broker-level effects. Because treated and control analysts are located within the same brokerage house around the same departure event, they are exposed to similar broker-level changes. Identification comes from comparing the change in forecast informativeness for analysts covering affected industries relative to analysts covering unaffected industries within the same broker and event window. The regression specification is identical to equation (1).

Figure 3, Panel A reports results for the stacked DiD estimates. In the 12-month window, the coefficient on $SignedRevision \times Treatment \times Post$ is -0.509. After a senior colleague departs, the market reaction to forecast revisions falls for remaining analysts covering the affected industry, relative to their colleagues covering other industries at the same broker. A one-standard-deviation increase in *SignedRevision* translates into a roughly 1 percentage-point smaller announcement-window CAR (-0.509 $\times$ 0.0194). The negative effect is consistent across windows, ranging from -0.765 at the 6-month horizon to -0.372 at the 21-month horizon, and attenuates as the window lengthens, consistent with teams gradually adjusting as the broker is reorganized or replacement expertise arrives.

Panel B indicates that the decline is unlikely to reflect a pre-existing differential trend. The pre-retirement coefficients at months -12, -9, and -6 are statistically indistinguishable from zero and exhibit no systematic downward pattern. The estimates become negative after the retirement, reaching -0.513 at month +3 and -0.675 at month +6, before attenuating to -0.512 at month +9 and -0.052 at month +12. This pattern suggests that the

[10] We focus on junior analysts because they stand to gain the most from a senior colleague's expertise, mentorship, and client access, so any complementarity within teams should be most detectable in this group. Restricting the sample to junior analysts also sharpens identification since it holds the experience composition of treated and control units fixed, ensuring that the estimated effect is not driven by differences in seniority between analysts exposed and unexposed to the departure.

[11] Analyst research departments are often structured along either country or industry lines (Sonney, 2009). Consistent with this institutional structure, Phua et al. (2023) show that analysts within the same brokerage house who cover common industry tend to share information. A broad industry classification is therefore useful for capturing potential team-level knowledge spillovers. Our results are quantitatively similar using the Fama-French 48 industry classification.

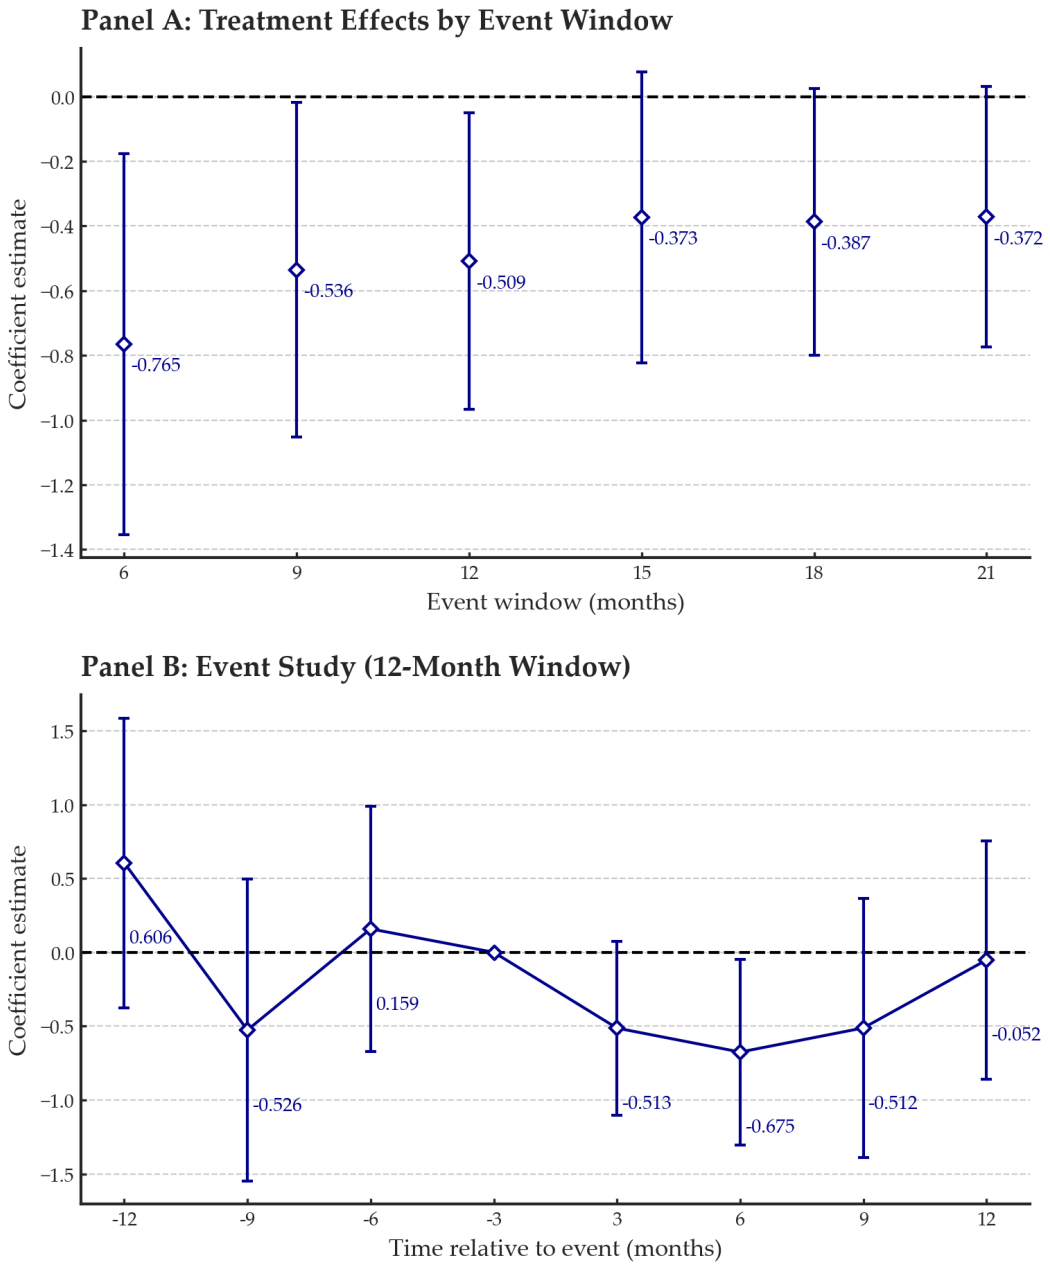


Figure 3. Senior Colleague Retirements and Analyst Forecast Informativeness.

Notes: This figure reports stacked DiD estimates of the effect of senior analyst retirements on the price informativeness of junior colleagues' forecast revisions. Treated analysts cover the retiring analyst's Fama-French 30 industry, while controls cover other industries within the same broker. The reported coefficient is on the triple interaction term *SignedRevision* × *Treatment* × *Post*. We infer analysts' ages and reconstruct their employment histories using their LinkedIn profiles. A senior retiree is an analyst who, at the time of departure, is older than 55 and has at least 20 years of financial-industry experience, permanently disappears from IBES, and is not subsequently observed moving to another firm or industry. We focus on remaining analysts with fewer than five years of experience. All specifications include analyst-event-firm and year-month fixed effects, with standard errors clustered by retirement event. Panel A reports the coefficient estimates from six separate regressions, each using a symmetric pre/post estimation window of the length shown on the horizontal axis. Panel B reports a dynamic event-study specification estimated on the 12-month window, plotting the triple-interaction coefficient by quarter relative to the retirement, with the quarter immediately preceding the retirement as the reference point. Point estimates are labeled, and vertical bars denote 95% confidence intervals.

loss of a senior colleague produces an immediate but temporary reduction in the informativeness of affected junior analysts' research. The absence of a systematic pre-trend weighs against selection on differential trends or anticipation of retirement. Furthermore, Figure A2 reports a placebo test using senior analysts who meet the same requirements but do not retire. The triple interaction is small and statistically insignificant across all specifications, supporting the interpretation that the effect arises only following an actual departure.

Taken together, this evidence indicates that analyst team composition matters for price informativeness. In contrast to the substitution pattern in the platform results, analysts of different types act as complements within teams. Senior analysts generate positive spillovers to their colleagues, either by providing industry expertise and internal advice or by drawing investor attention and client engagement to the broker's research in the affected industry.

2.5 Analyst Job Mobility Is Declining

So far, sections 2.3 and 2.4 have shown, respectively, the importance of analyst-broker matching and analyst-team composition for price efficiency. However, various labor-market frictions may impede the efficient allocation of analysts across brokers, with imperfect job mobility being one of the most prominent examples.

When job switching entails high costs, such as litigation risk, reputational damage, and network losses, better-matched alternative brokers must offer analysts sufficiently high compensation to attract them. When the required compensation exceeds the surplus generated by the analyst, the better match fails to materialize. From the analyst's perspective, job mobility is closely related to the option value of future job switching. If moving away from a currently well-matched broker in the future involves high job-switching costs, career concerns may prevent the analyst from choosing that broker in the first place.

Panel A of Figure 4 plots two time series of annual analyst job-switching rates across brokers from 1995 to 2024, one using IBES data and the other using LinkedIn data.[12] Two

[12] For each analyst-broker-year, we count the number of forecasts issued by the analyst at the broker, the number of distinct firms covered, and the date of the analyst's last forecast. We then assign each analyst to a primary broker in each year, defined as the broker for which the analyst issues the largest number of forecasts. Ties are broken first by the number of firms covered and then by the last forecast

patterns stand out. First, the labor market for financial analysts features low job mobility. In the IBES data, the average annual job-switching rate over the sample period is 8.89%, with a peak of 14.56% in 2001. By comparison, economy-wide survey evidence shows that the overall annual employer-to-employer transition rate over the same period is about 25.01% (Fujita et al., 2024).[13] Second, analyst job-switching rates have declined persistently since the financial crisis, falling to only 4.06% in 2024 based on the IBES data. The LinkedIn data show the same patterns. Extremely low job mobility implies that high-quality analyst-broker matches are harder to form, thereby impeding price discovery.

Panel B of Figure 4 focuses on the post-2010 period during which analyst job mobility has been declining, and plots average job-switching rates by analyst experience.[14] For analysts with less than 20 years of experience in the IBES data, job-switching rates are relatively stable, with an average of 7.07%. After 20 years of experience, however, job mobility starts to decline rapidly, and the average job-switching rate is only 4.11%. For the LinkedIn data, perhaps because analysts who actively pursue job changes selectively enter the LinkedIn sample, job-switching rates decrease almost monotonically with analyst experience. Compared with less experienced analysts, senior analysts may carry more client relationships and organizational capital for brokers, and may be subject to stricter contractual terms and stronger reputational concerns, making job changes even more costly. However, as shown in section 2.4, these experienced analysts often play a critical role within teams.

2.6 Discussion

Taken together, our motivating evidence suggests that although brokerage platforms and analyst-team composition are both crucial for the informativeness of analyst research, declining analyst job mobility, especially among senior analysts who are central to team pro-

date. We restrict attention to analyst-year observations that can be linked to the same analyst in the previous calendar year. An analyst is defined as switching employers in year y if the primary broker in year y differs from that in year $y-1$. The annual job-switching rate is the fraction of such analyst-year transitions in which the analyst changes primary brokers.

[13] Fujita et al. (2024) report an average monthly employer-to-employer transition rate of 2.37% from October 1995 to December 2024, which implies an average annual transition rate of $1-(1-0.0237)^{12} = 0.2501$.

[14] We measure experience by the analyst's tenure at the origin broker before the job switch is realized. We round analyst tenure to the nearest year and compute the average job switching rate within each tenure bin, requiring at least 30 analyst transitions for each bin.

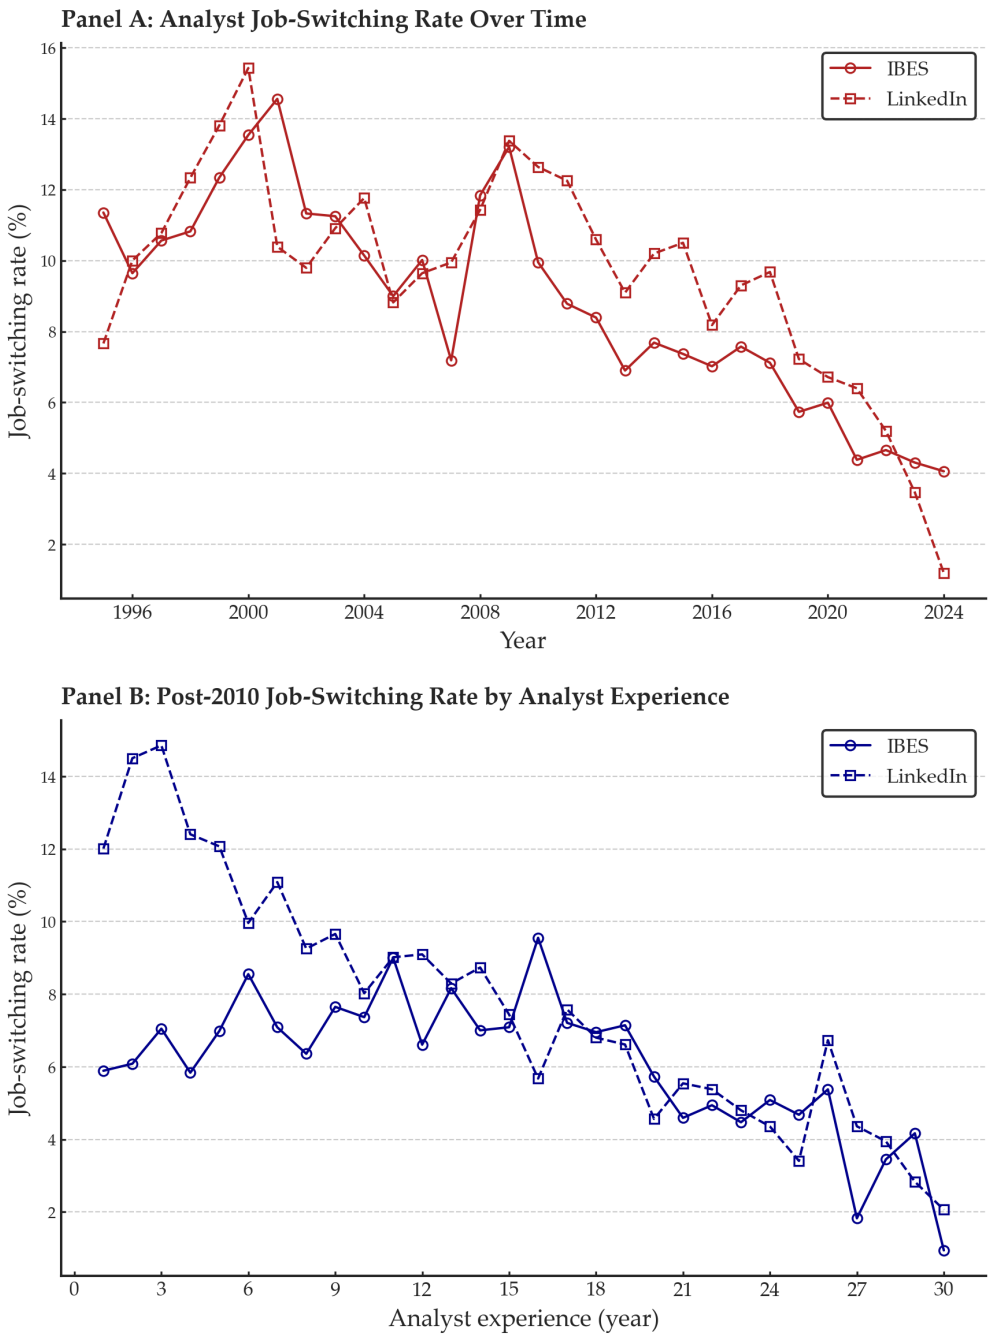


Figure 4. Analyst Job Mobility.

Notes: Panel A plots the time series of annual analyst job-switching rates across brokers from 1995 to 2024. Panel B focuses on the post-2010 period and plots average job-switching rates by analyst experience. The solid and dashed lines refer to the IBES and LinkedIn data, respectively.

duction, may distort analyst-broker matching and ultimately reduce analysts' aggregate contribution to price efficiency. To unpack the role of different labor-market mechanisms in allocating analysts across brokerage platforms and quantify the aggregate impact of this allocation on stock price efficiency, we develop a structural model centered on the labor market for financial analysts in the next section.

In the model, brokers compete for analyst labor in an oligopolistic labor market, using high-quality information production to attract client demand and extract surplus. We allow for flexible matching mechanisms between brokers and analyst teams, as well as among different analyst types within teams. Analysts make dynamic forward-looking broker choices, taking into account wage flows, future job-switching opportunities, mobility costs, and idiosyncratic preference shocks. Equilibrium wage payoffs are determined by Nash bargaining between brokers and analysts within a Nash equilibrium, taking both parties' surpluses into account. Together, these components endogenously determine the allocation of analysts across brokers, which in turn shapes the quality of analyst research and the aggregate impact on price efficiency. Figure 1 illustrates the structure of the model.

3 Model

3.1 Environment

Client preferences There is a unit mass of clients of brokerage services, indexed by $o \in [0, 1]$. There is a finite number N of brokers, indexed by $i \in \{1, \ldots, N\}$. A given client o chooses the broker i that maximizes utility from brokerage services:

$$U(o) = \max_i \left\{ f_i^{-\kappa} \varepsilon_i(o) \right\},$$

where f_i is the fee endogenously set by broker i, $\kappa > 0$ denotes the elasticity of utility with respect to brokerage fee, and $\varepsilon_i(o)$ represents client o's idiosyncratic preference for broker i.

We assume that idiosyncratic preferences for a given broker i are i.i.d. across clients

and follow a Fréchet distribution with cumulative distribution function (CDF) given by:[15]

$$G_i(\varepsilon) = \exp\left(-q_i \varepsilon^{-\theta}\right),$$

where q_i denotes broker i's composite quality, which is endogenously determined by the bargaining and matching process between brokers and analysts. A higher q_i implies that the realization of idiosyncratic preferences for broker i is higher on average, indicating a higher absolute advantage of broker i in attracting clients. We will specify this term later. The term θ is the shape parameter of the Fréchet distribution, with a lower θ indicating greater dispersion of idiosyncratic preferences across clients.

Pricing Brokers engage in price competition à la Bertrand, setting fees to maximize total surplus while taking other brokers' fees as given. Let λ_i denote the fraction of clients choosing broker i. The objective function is given by:

$$\max_{f_i} \lambda_i (f_i - c), \quad \forall\, i \in \{1, \ldots, N\},$$

where c denotes the marginal cost, which we assume to be constant across brokers. Solving the optimal pricing problem yields the following proposition:

Proposition 1 (Total Surplus of Brokers)**.** *Let $\tilde{\theta} = \kappa\theta > 1$. The optimal fee charged by broker i satisfies $f_i = m(\lambda_i)\, c$, where $m(\lambda_i) = \frac{\sigma(\lambda_i)}{\sigma(\lambda_i)-1}$ denotes the markup and $\sigma(\lambda_i) = \tilde{\theta}(1-\lambda_i)$ represents the demand elasticity. Client shares $\{\lambda_i\}_{i=1}^N$ are, in turn, determined by the following system of equations:*

$$\lambda_i = \frac{q_i m(\lambda_i)^{-\tilde{\theta}}}{\sum_{i'=1}^N q_{i'} m(\lambda_{i'})^{-\tilde{\theta}}}, \quad \forall\, i \in \{1, \ldots, N\}. \tag{2}$$

Given an $N \times 1$ quality vector $\boldsymbol{q} = [q_1, \ldots, q_N]^\top$, the total surplus of broker i is given by:

$$S_i^T(\boldsymbol{q}) = \frac{\lambda_i(\boldsymbol{q})}{\tilde{\theta}[1 - \lambda_i(\boldsymbol{q})] - 1} c, \tag{3}$$

where $\{\lambda_i(\boldsymbol{q})\}_{i=1}^N$ solves (2).

[15] Assuming Fréchet shocks is standard in settings with a max operator because the Fréchet distribution is max-stable and delivers tractable closed-form choice probabilities (Lind and Ramondo, 2023). This assumption has been widely used across multiple fields since the Ricardian trade model of Eaton and Kortum (2002).

Proof. See Appendix B.1. □

The optimal fee f_i resembles standard oligopolistic pricing (Atkeson and Burstein, 2008), in that the variable markup $m(\lambda_i)$ charged by broker i is decreasing in the demand elasticity it faces, and that the demand elasticity $\sigma(\lambda_i)$ is decreasing in client share λ_i. Intuitively, a larger client base implies greater market power, making demand less price-sensitive and allowing the broker to charge a higher markup. The demand elasticity also depends on the composite parameter $\tilde{\theta}$. Demand is more elastic to fee changes when clients' utility is more sensitive to the service fee (i.e., higher κ) and when clients' idiosyncratic preferences are less dispersed (i.e., higher θ).

Composite quality There is a unit mass of atomistic analysts. The composite quality term q_i aggregates an analyst-team-specific component and a platform-specific component:

$$q_i = \left(\Omega_i^{\eta} + b_i^{\eta}\right)^{\frac{1}{\eta}}, \quad \Omega_i = \left[\sum_{x=1}^{\overline{x}} \left(a_{ix}\omega_x\right)^{\gamma}\right]^{\frac{1}{\gamma}}, \tag{4}$$

where Ω_i denotes the quality of the analyst team in broker i, with ω_x representing the quality of type-x analysts and a_{ix} being the mass of type-x analysts in broker i. We assume that there is a finite number $\overline{x}$ of analyst quality types and index them by $x \in \{1, \ldots, \overline{x}\}$. The term b_i denotes the platform quality component. Parameter η governs the interaction between analyst team quality and platform quality: $0 < \eta < 1$ implies complementarity, $\eta > 1$ implies substitutability, and $\eta = 1$ implies independence.[16] Parameter γ governs the within-team interaction across analyst types: $0 < \gamma < 1$ implies complementarity, $\gamma > 1$ implies substitutability, and $\gamma = 1$ implies independence.[17]

To facilitate representation and computation, let $\boldsymbol{a}_i = [a_{i1}, \ldots, a_{i\overline{x}}]^{\top}$ denote an $\overline{x} \times 1$ analyst mass vector for broker i, let $\boldsymbol{\omega} = [\omega_1, \ldots, \omega_{\overline{x}}]^{\top}$ denote an $\overline{x} \times 1$ analyst-type-specific quality vector, and let $\tilde{\boldsymbol{b}} = [b_1^{\eta}, \ldots, b_N^{\eta}]^{\top}$ denote an $N \times 1$ effective platform quality vector.

[16] The cross-partial is given by:

$$\frac{\partial^2 q_i}{\partial \Omega_i \partial b_i} = (1-\eta)\left(\Omega_i b_i\right)^{\eta-1} q_i^{1-2\eta}.$$

[17] For $x \neq x'$, the cross-partial is given by:

$$\frac{\partial^2 \Omega_i}{\partial a_{ix} \partial a_{ix'}} = (1-\gamma)\left(a_{ix}a_{ix'}\right)^{\gamma-1}\left(\omega_x \omega_{x'}\right)^{\gamma} \Omega_i^{1-2\gamma}.$$

All analyst-type- and platform-specific quality components, i.e., $\boldsymbol{\omega}$ and $\tilde{\boldsymbol{b}}$, are treated as exogenous endowments. Analyst allocation across brokers, i.e., $\boldsymbol{a}_i$ for all $i \in \{1, \ldots, N\}$, is endogenously determined by the bargaining and matching process between brokers and analysts, which we specify later.

Analyst preferences Time is discrete and indexed by t. The lifetime value function of a representative type-x analyst in broker i at time t is characterized by the following Bellman equation:

$$v_{ixt} = \ln w_{ixt} + \max_{i'} \left\{ \beta \mathbb{E} \left[v_{i'xt+1} \right] - \tau_{ii'} + \nu \varphi_{i'xt} \right\}.$$

At time t, the analyst first receives the flow log utility from payoff w_{ixt}. We assume that the payoff offered by broker $i \in \{1, \ldots, N\}$ at time t is preset by a bargaining process at time $t-1$, which we specify later. We also allow for a composite outside option indexed by $i = 0$, with an exogenous payoff $w_{0xt} = \overline{w}_{xt}$.

Next, the type-x analyst realizes the idiosyncratic preferences for each choice $i' \in \{0, 1, \ldots, N\}$, denoted by $\varphi_{i'xt}$, and then chooses the option i' that maximizes the net discounted expected value, $\beta \mathbb{E} \left[v_{i'xt+1} \right] - \tau_{ii'} + \nu \varphi_{i'xt}$, where β denotes the discount factor, $\nu > 1$ is a scale parameter that controls the variance of idiosyncratic preference shocks, and $\tau_{ii'}$ represents the job-switching cost from i to i'.[18] We assume $\tau_{ii'} = 0$ for $i' = i$ and $\tau_{ii'} > 0$ for $i' \neq i$. Finally, we assume that idiosyncratic preferences are i.i.d. across analysts and options and follow a Gumbel distribution with CDF given by:

$$F(\varphi) = \exp\left(-\exp\left(-\varphi - \overline{\gamma} \right) \right),$$

where $\overline{\gamma} = -\Gamma'(1) \approx 0.57721$ is the Euler's constant.[19]

Let $V_{ixt} = \mathbb{E}\left[v_{ixt} \right]$ denote the expected value before the realization of idiosyncratic preference shocks, we have:

$$V_{ixt} = \ln w_{ixt} + \mathbb{E}\left[\max_{i'} \left\{ \beta V_{i'xt+1} - \tau_{ii'} + \nu \varphi_{i'xt} \right\} \right].$$

For tractability in the bargaining process, we focus on the steady state in which all vari-

[18] In reality, the job-switching cost includes search and application costs, administrative and legal frictions, network costs, etc.

[19] Assuming Gumbel-distributed idiosyncratic preferences has been widely used since Caliendo et al. (2019) to characterize labor market dynamics.

ables remain constant over time. Taking payoffs as given, solving the expected value maximization problem of analysts yields the following proposition:

Proposition 2 (Steady-State Analyst Allocation)**.** *In steady state, the money-metric expected value for a representative type-x analyst in broker i is given by:*

$$\widetilde{V}_{ix} = w_{ix}\left[\sum_{i'=0}^{N}\exp\left(-\frac{\tau_{ii'}}{\nu}\right)\widetilde{V}_{i'x}^{\frac{\beta}{\nu}}\right]^{\nu}, \quad \forall\, i \in \{0,1,\ldots,N\}\,. \tag{5}$$

Given an $(N+1)\times 1$ payoff vector $\boldsymbol{w}_x = [w_{0x}, w_{1x}, \ldots, w_{Nx}]^\top$, the mass of type-$x$ analysts in broker i is given by:

$$a_{ix} = \sum_{i'=0}^{N}\mu_{i'ix}\left(\boldsymbol{w}_x\right)a_{i'x}, \quad \forall\, i \in \{0,1,\ldots,N\}\,, \tag{6}$$

where

$$\mu_{i'ix}\left(\boldsymbol{w}_x\right) = \frac{\exp\left(-\frac{\tau_{i'i}}{\nu}\right)\widetilde{V}_{ix}\left(\boldsymbol{w}_x\right)^{\frac{\beta}{\nu}}}{\sum_{k=0}^{N}\exp\left(-\frac{\tau_{i'k}}{\nu}\right)\widetilde{V}_{kx}\left(\boldsymbol{w}_x\right)^{\frac{\beta}{\nu}}} \tag{7}$$

is the fraction of type-x analysts switching jobs from i' to i and $\left\{\widetilde{V}_{ix}\left(\boldsymbol{w}_x\right)\right\}_{i=0}^{N}$ solves (5). The solution to (6) is denoted by $\left\{a_{ix}\left(\boldsymbol{w}_x\right)\right\}_{i=0}^{N}$.

Proof. See Appendix B.2. □

Bargaining For type-x analysts, the elements $\{w_{ix}\}_{i=1}^{N}$ of the payoff vector $\boldsymbol{w}_x$ are determined through bargaining with brokers $i \in \{1,\ldots,N\}$. For tractability, we impose the following assumption on the bargaining process:[20]

Assumption 1 (Bargaining Process)**.** *The following restrictions apply:*

- ***Bilateral bargaining.*** *Bargaining occurs bilaterally between an infinitesimal mass $\epsilon \to 0$ of analysts and the current employer broker. Collective bargaining by a set of analysts and multilateral bargaining with other brokers are not allowed.*
- ***No commitment.*** *Payoffs are renegotiated one period ahead, before the realization of analysts' idiosyncratic preferences, and are always rejectable. Long-term contracts that commit to a contingent payoff path are not allowed.*

[20] Similar assumptions are made by Eaton et al. (2026) in the context of buyer-supplier bargaining in international markets.

- ***Simultaneity.*** *All analyst-broker negotiations over next-period payoffs are resolved simultaneously. No negotiation is conditioned on the realized outcome of any other negotiation within the same period.*

- ***No discrimination.*** *Each broker i can offer only a single payoff w_{ix} that applies to all type-x analysts it employs in the next period, including those who switch from another broker $i' \neq i$ to i.*

- ***Passive beliefs.*** *In any bargaining between a type-x analyst and a broker i, the two parties take as given the equilibrium payoffs set by all other brokers $i' \neq i$, denoted by the $(N-1)\times 1$ vector $\boldsymbol{w}^*_{\neg ix}$. Any off-equilibrium action by either party does not change their beliefs about $\boldsymbol{w}^*_{\neg ix}$.*

Under Assumption 1, bargaining between brokers and analysts takes the form of Nash bargaining embedded in the Nash equilibrium of a non-cooperative game. Specifically, in each broker i, the type-x analyst payoff w_{ix} is set to maximize the exponentiated product of the two parties' surpluses from agreement relative to disagreement, taking as given the equilibrium payoffs set by all other brokers, $\boldsymbol{w}^*_{\neg ix}$. Each broker i and each type-x analyst therefore treat the money-metric expected value $\widetilde{V}_{ix}$ and the analyst mass a_{ix} as functions only of the payoff w_{ix}, denoted by $\widetilde{V}_{ix}\left(w_{ix};\boldsymbol{w}^*_{\neg ix}\right)$ and $a_{ix}\left(w_{ix};\boldsymbol{w}^*_{\neg ix}\right)$, respectively. Solving the bargaining problem yields the following proposition:

Proposition 3 (Equilibrium Payoffs)**.** *Let α_x denote the bargaining power of type-x analysts, $\frac{\partial f(\boldsymbol{x})}{\partial x}\big|_{\boldsymbol{x}=\boldsymbol{x}^*} = \frac{\partial f(\boldsymbol{x}^*)}{\partial x^*}$ denote the partial derivative of function $f(\boldsymbol{x})$ with respect to x evaluated at the equilibrium value $\boldsymbol{x}=\boldsymbol{x}^*$, and $\boldsymbol{W}^* = [\boldsymbol{w}^*_1,\ldots,\boldsymbol{w}^*_{\overline{x}}]$ denote an $(N+1)\times\overline{x}$ equilibrium payoff matrix. The equilibrium payoff set by broker i for type-x analysts is given by:*

$$w^*_{ix} = \widetilde{V}_{ix}\left(\boldsymbol{w}^*_x\right)^{\frac{1-\beta}{1-\nu}}\left[\frac{\partial\widetilde{V}_{ix}\left(\boldsymbol{w}^*_x\right)}{\partial w^*_{ix}}\right]^{\frac{\nu}{\nu-1}}\left\{\frac{(1-\alpha_x)\,S^A_{ix}\left(\boldsymbol{w}^*_x\right)}{\beta\alpha_x S^B_{ix}\left(\boldsymbol{W}^*\right)}\left[1-\frac{\partial}{\partial w^*_{ix}}\left(\frac{\partial S^T_i\left(\boldsymbol{q}\left(\boldsymbol{W}^*\right)\right)}{\partial a_{ix}\left(\boldsymbol{w}^*_x\right)}\right)\right]\right\}^{\frac{\nu}{1-\nu}}, \tag{8}$$

where

$$S^A_{ix}\left(\boldsymbol{w}^*_x\right) = \left[\sum_{i'=0}^{N}\exp\left(-\frac{\tau_{ii'}}{\nu}\right)\widetilde{V}_{i'x}\left(\boldsymbol{w}^*_x\right)^{\frac{\beta}{\nu}}\right]^{\nu} - \left[\sum_{i'\neq i}\exp\left(-\frac{\tau_{ii'}}{\nu}\right)\widetilde{V}_{i'x}\left(\boldsymbol{w}^*_x\right)^{\frac{\beta}{\nu}}\right]^{\nu} \tag{9}$$

captures type-x analysts' surplus from broker i in the Nash bargaining objective, and

$$S_{ix}^{B}\left(\boldsymbol{W}^{*}\right)=\frac{\partial S_{i}^{T}\left(\boldsymbol{q}\left(\boldsymbol{W}^{*}\right)\right)}{\partial a_{ix}\left(\boldsymbol{w}_{x}^{*}\right)}-\sum_{i'\neq i}\frac{\exp\left(-\frac{\tau_{ii'}}{\nu}\right)\widetilde{V}_{i'x}\left(\boldsymbol{w}_{x}^{*}\right)^{\frac{\beta}{\nu}}}{\sum_{k\neq i}\exp\left(-\frac{\tau_{ik}}{\nu}\right)\widetilde{V}_{kx}\left(\boldsymbol{w}_{x}^{*}\right)^{\frac{\beta}{\nu}}}\frac{\partial S_{i}^{T}\left(\boldsymbol{q}\left(\boldsymbol{W}^{*}\right)\right)}{\partial a_{i'x}\left(\boldsymbol{w}_{x}^{*}\right)}-w_{ix}^{*} \tag{10}$$

captures broker i's surplus from type-x analysts in the Nash bargaining objective.

Proof. See Appendix B.3. □

3.2 Steady-State Equilibrium

To close the model, we require the following labor market clearing conditions:

$$\sum_{i=0}^{N}a_{ix}=\overline{a}_{x},\ \ \forall\, x\in\{1,\ldots,\overline{x}\}\,, \tag{11}$$

where $\overline{a}_x$ denotes the exogenous total endowment of type-x analysts, and

$$\sum_{x=1}^{\overline{x}}\overline{a}_{x}=\overline{\boldsymbol{a}}^{\top}\mathbf{1}=1. \tag{12}$$

Finally, define an $\overline{x}\times 1$ outside payoff vector $\overline{\boldsymbol{w}}=[\overline{w}_1,\ldots,\overline{w}_{\overline{x}}]^{\top}$, an $(N+1)\times(N+1)$ inverse job-switching cost matrix $\widetilde{\boldsymbol{\tau}}$ with (i,i') entry $\exp\left(-\frac{\tau_{ii'}}{\nu}\right)$, and an $\overline{x}\times 1$ analyst bargaining power vector $\boldsymbol{\alpha}=[\alpha_1,\ldots,\alpha_{\overline{x}}]^{\top}$, we characterize the steady-state equilibrium as follows:

Definition 1 (Steady-State Equilibrium)**.** *Given the set of parameters:*

$$\boldsymbol{\Theta}=\left\{N,\overline{x},c,\tilde{\theta},\eta,\gamma,\beta,\nu,\boldsymbol{\omega},\tilde{\boldsymbol{b}},\overline{\boldsymbol{w}},\widetilde{\boldsymbol{\tau}},\boldsymbol{\alpha},\overline{\boldsymbol{a}}\right\},$$

the steady-state equilibrium is characterized by an equilibrium payoff matrix $\boldsymbol{W}^{*}$ such that:

- *Analysts of all types $x\in\{1,\ldots,\overline{x}\}$ currently in $i\in\{0,1,\ldots,N\}$ choose the next-period job option $i'\in\{0,1,\ldots,N\}$ to maximize their lifetime value. The job-switching fraction $\mu_{i'ix}\left(\boldsymbol{w}_{x}^{*}\right)$ is characterized by equation (7), in which the money-metric expected value before the realization of idiosyncratic preferences, $\widetilde{V}_{ix}\left(\boldsymbol{w}_{x}\right)$, solves the system of equations (5).*

- *The allocation of all types $x \in \{1, \ldots, \overline{x}\}$ of analysts across job options $i' \in \{0, 1, \ldots, N\}$ is summarized by an $(N+1) \times \overline{x}$ matrix $\boldsymbol{A}(\boldsymbol{W}^*) = [\boldsymbol{a}_0(\boldsymbol{W}^*), \ldots, \boldsymbol{a}_N(\boldsymbol{W}^*)]^\top$, in which each element $a_{ix}(\boldsymbol{w}_x^*)$ solves the system of equations (6) together with the analyst labor market clearing conditions (11) and (12).*

- *Each broker $i \in \{1, \ldots, N\}$ sets service fee to maximize its total surplus earned from clients. The total surplus $S_i^T(\boldsymbol{q}(\boldsymbol{W}^*))$ is given by equation (3), in which the client share $\lambda_i(\boldsymbol{q}(\boldsymbol{W}^*))$ solves the system of equations (2), and each element in the quality vector $\boldsymbol{q}(\boldsymbol{W}^*) = [q_1(\boldsymbol{a}_1(\boldsymbol{W}^*)), \ldots, q_N(\boldsymbol{a}_N(\boldsymbol{W}^*))]^\top$ is given by equation (4).*

- *The equilibrium payoff matrix $\boldsymbol{W}^*$, in turn, is determined by Nash bargaining between analysts and brokers under Assumption 1, based on the two parties' relevant surpluses given by (9) and (10), respectively. In particular, each element of $\boldsymbol{W}^*$, except those in the exogenously given first row, solves the system of equations (8).*

3.3 Solution Algorithm

To solve for the steady-state equilibrium, it remains to obtain the gradient terms in equations (8) and (10). The following proposition derives these gradients step by step:

Proposition 4 (Gradients)**.** *The marginal effect of the payoff from option i on type-x analysts' money-metric expected value in option k is given by:*

$$\boxed{\frac{\partial \widetilde{V}_{kx}(\boldsymbol{w}_x)}{\partial w_{ix}}} = \frac{\widetilde{V}_{kx}(\boldsymbol{w}_x)}{w_{ix}} \cdot \frac{\partial \ln \widetilde{V}_{kx}(\boldsymbol{w}_x)}{\partial \ln w_{ix}}, \tag{13}$$

where the elasticities $\frac{\partial \ln \widetilde{V}_{kx}(\boldsymbol{w}_x)}{\partial \ln w_{ix}}$ for all $k \in \{0, 1, \ldots, N\}$ solves the following system of equations:

$$\frac{\partial \ln \widetilde{V}_{kx}(\boldsymbol{w}_x)}{\partial \ln w_{ix}} = \mathbf{1}\{k = i\} + \beta \sum_{i'=0}^{N} \mu_{ki'x}(\boldsymbol{w}_x) \frac{\partial \ln \widetilde{V}_{i'x}(\boldsymbol{w}_x)}{\partial \ln w_{ix}}. \tag{14}$$

The marginal effect of type-x analysts in broker k on broker i's total surplus is given by:

$$\boxed{\frac{\partial S_i^T(\boldsymbol{q}(\boldsymbol{W}))}{\partial a_{kx}(w_{kx}; \boldsymbol{w}_{\neg kx})}} = \frac{\partial S_i^T(\boldsymbol{q}(\boldsymbol{W}))}{\partial \lambda_i(\boldsymbol{q}(\boldsymbol{W}))} \cdot \frac{\partial \lambda_i(\boldsymbol{q}(\boldsymbol{W}))}{\partial q_k(\boldsymbol{a}_k(\boldsymbol{W}))} \cdot \frac{\partial q_k(\boldsymbol{a}_k(\boldsymbol{W}))}{\partial a_{kx}(w_{kx}; \boldsymbol{w}_{\neg kx})}, \tag{15}$$

where

$$\frac{\partial S_i^T\left(\boldsymbol{q}\left(\boldsymbol{W}\right)\right)}{\partial \lambda_i\left(\boldsymbol{q}\left(\boldsymbol{W}\right)\right)} = \frac{\left(\tilde{\theta}-1\right)c}{\left[\tilde{\theta}\left(1-\lambda_i\left(\boldsymbol{q}\left(\boldsymbol{W}\right)\right)\right)-1\right]^2} \tag{16}$$

is the marginal effect of broker i's client share on its total surplus,

$$\frac{\partial \lambda_i\left(\boldsymbol{q}\left(\boldsymbol{W}\right)\right)}{\partial q_k\left(\boldsymbol{a}_k\left(\boldsymbol{W}\right)\right)} = \frac{\lambda_i\left(\boldsymbol{q}\left(\boldsymbol{W}\right)\right)}{q_k\left(\boldsymbol{a}_k\left(\boldsymbol{W}\right)\right)} \cdot \frac{\partial \ln \lambda_i\left(\boldsymbol{q}\left(\boldsymbol{W}\right)\right)}{\partial \ln q_k\left(\boldsymbol{a}_k\left(\boldsymbol{W}\right)\right)} \tag{17}$$

is the marginal effect of broker k's composite quality on broker i's client share, with the elasticities $\frac{\partial \ln \lambda_i(\boldsymbol{q}(\boldsymbol{W}))}{\partial \ln q_k(\boldsymbol{a}_k(\boldsymbol{W}))}$ for all $i \in \{1,\ldots,N\}$ solving the following system of equations:

$$\begin{aligned}\frac{\partial \ln \lambda_i\left(\boldsymbol{q}\left(\boldsymbol{W}\right)\right)}{\partial \ln q_k\left(\boldsymbol{a}_k\left(\boldsymbol{W}\right)\right)} = {} & \left[\mathbf{1}\left\{i=k\right\} - \lambda_k\left(\boldsymbol{q}\left(\boldsymbol{W}\right)\right)\right] \\ & - \tilde{\theta}\sum_{i'=1}^{N} \frac{\lambda_{i'}\left(\boldsymbol{q}\left(\boldsymbol{W}\right)\right)\left[\mathbf{1}\left\{i'=i\right\} - \lambda_{i'}\left(\boldsymbol{q}\left(\boldsymbol{W}\right)\right)\right]}{\left(1-\lambda_{i'}\left(\boldsymbol{q}\left(\boldsymbol{W}\right)\right)\right)\left[\tilde{\theta}\left(1-\lambda_{i'}\left(\boldsymbol{q}\left(\boldsymbol{W}\right)\right)\right)-1\right]} \cdot \frac{\partial \ln \lambda_{i'}\left(\boldsymbol{q}\left(\boldsymbol{W}\right)\right)}{\partial \ln q_k\left(\boldsymbol{a}_k\left(\boldsymbol{W}\right)\right)},\end{aligned} \tag{18}$$

and

$$\frac{\partial q_k\left(\boldsymbol{a}_k\left(\boldsymbol{W}\right)\right)}{\partial a_{kx}\left(w_{kx};\boldsymbol{w}_{\neg kx}\right)} = q_k\left(\boldsymbol{a}_k\left(\boldsymbol{W}\right)\right)^{1-\eta}\Omega_k\left(\boldsymbol{a}_k\left(\boldsymbol{W}\right)\right)^{\eta-\gamma} a_{kx}\left(w_{kx};\boldsymbol{w}_{\neg kx}\right)^{\gamma-1}\omega_x^{\gamma} \tag{19}$$

is the marginal effect of type-x analysts on broker k's composite quality. The payoff sensitivity of broker i's marginal total surplus from its own type-x analysts is given by:

$$\boxed{\frac{\partial}{\partial w_{ix}}\left(\frac{\partial S_i^T\left(\boldsymbol{q}\left(\boldsymbol{W}\right)\right)}{\partial a_{ix}\left(w_{ix};\boldsymbol{w}_{\neg ix}\right)}\right)} = \frac{\partial S_i^T\left(\boldsymbol{q}\left(\boldsymbol{W}\right)\right)}{\partial \lambda_i\left(\boldsymbol{q}\left(\boldsymbol{W}\right)\right)} \cdot \frac{\partial \lambda_i\left(\boldsymbol{q}\left(\boldsymbol{W}\right)\right)}{\partial q_i\left(\boldsymbol{a}_i\left(\boldsymbol{W}\right)\right)} \cdot \frac{\partial}{\partial w_{ix}}\left(\frac{\partial q_i\left(\boldsymbol{a}_i\left(\boldsymbol{W}\right)\right)}{\partial a_{ix}\left(w_{ix};\boldsymbol{w}_{\neg ix}\right)}\right), \tag{20}$$

where

$$\begin{aligned}\frac{\partial}{\partial w_{ix}}\left(\frac{\partial q_i\left(\boldsymbol{a}_i\left(\boldsymbol{W}\right)\right)}{\partial a_{ix}\left(w_{ix};\boldsymbol{w}_{\neg ix}\right)}\right) = {} & \left[\frac{1-\eta}{q_i\left(\boldsymbol{a}_i\left(\boldsymbol{W}\right)\right)} \cdot \frac{\partial q_i\left(\boldsymbol{a}_i\left(\boldsymbol{W}\right)\right)}{\partial a_{ix}\left(w_{ix};\boldsymbol{w}_{\neg ix}\right)} + \frac{\left(\eta-\gamma\right)a_{ix}\left(w_{ix};\boldsymbol{w}_{\neg ix}\right)^{\gamma-1}\omega_x^{\gamma}}{\Omega_i\left(\boldsymbol{a}_i\left(\boldsymbol{W}\right)\right)^{\gamma}}\right. \\ & \left. + \frac{\gamma-1}{a_{ix}\left(w_{ix};\boldsymbol{w}_{\neg ix}\right)}\right] \cdot \frac{\partial q_i\left(\boldsymbol{a}_i\left(\boldsymbol{W}\right)\right)}{\partial a_{ix}\left(w_{ix};\boldsymbol{w}_{\neg ix}\right)} \cdot \frac{\partial a_{ix}\left(w_{ix};\boldsymbol{w}_{\neg ix}\right)}{\partial w_{ix}}\end{aligned} \tag{21}$$

is the payoff sensitivity of broker i's marginal composite quality from its own type-x analysts, with the marginal effect of the payoff from option i on the allocation of type-x analysts across all options,

$\frac{\partial a_{kx}(\boldsymbol{w}_x)}{\partial w_{ix}}$ *for all* $k \in \{0, 1, \ldots, N\}$*, solving the following system of equations:*

$$\frac{\partial a_{kx}\left(\boldsymbol{w}_x\right)}{\partial w_{ix}} = \sum_{i'=0}^{N} \left[\frac{\mu_{i'kx}\left(\boldsymbol{w}_x\right)}{w_{ix}} \cdot \frac{\partial \ln \mu_{i'kx}\left(\boldsymbol{w}_x\right)}{\partial \ln w_{ix}} a_{i'x}\left(\boldsymbol{w}_x\right) + \mu_{i'kx}\left(\boldsymbol{w}_x\right) \frac{\partial a_{i'x}\left(\boldsymbol{w}_x\right)}{\partial w_{ix}} \right], \tag{22}$$

in which the elasticities of type-x *analysts' job-switching fractions from option* i' *to each option* $k \in \{0, 1, \ldots, N\}$ *with respect to the payoff from option* i *solve the following system of equations:*

$$\frac{\partial \ln \mu_{i'kx}\left(\boldsymbol{w}_x\right)}{\partial \ln w_{ix}} = \frac{1}{\nu} \left[\mathbf{1}\left\{i' = i\right\} + \beta \frac{\partial \ln \widetilde{V}_{kx}\left(\boldsymbol{w}_x\right)}{\partial \ln w_{ix}} - \frac{\partial \ln \widetilde{V}_{i'x}\left(\boldsymbol{w}_x\right)}{\partial \ln w_{ix}} \right]. \tag{23}$$

Proposition 4 provides a comprehensive picture of the model's inner workings. All endogenous variables can be expressed as functions of analysts' payoffs, which are in turn endogenously determined by Nash bargaining in a Nash equilibrium. Shocks to the parameter set $\boldsymbol{\Theta}$ imply changes in the equilibrium payoff matrix $\boldsymbol{W}^*$ and thus in all endogenous variables.

From equation (16), the assumption $\tilde{\theta} > 1$ ensures that the total surplus of broker i is strictly increasing in its own client share, i.e., $\frac{\partial S_i^T}{\partial \lambda_i} > 0$. Equation (18) implies that an increase in composite quality $\ln q_k$ affects client share $\ln \lambda_i$ through two channels. First, higher quality *directly* attracts more clients to broker k, as characterized by the first term in equation (18). Clearly, this term is strictly positive for $i = k$ and strictly negative for $i \neq k$. Second, changes in quality also *indirectly* affect client shares by shifting optimal markups, as captured by the second term in equation (18). As broker k's quality affects client shares and markups of all brokers, the second term is a function of $\left\{\frac{\partial \ln \lambda_i}{\partial \ln q_k}\right\}_{i=1}^{N}$, making (18) a system of equations for all $i \in \{1, \ldots, N\}$.

From equation (19), the marginal effect of type-x analysts on broker k's composite quality depends on the composition of *all* types of analysts $\boldsymbol{a}_k$, platform quality b_k, and parameters η and γ. Taking logs of equation (19) yields the following decomposition:

$$\ln\left(\frac{\partial q_k}{\partial a_{kx}}\right) = (1-\eta)\ln q_k + (\eta-\gamma)\ln \Omega_k + (\gamma-1)\ln a_{kx} + \gamma \ln \omega_x. \tag{24}$$

The marginal contribution of type-x analysts to broker k's composite quality can be decomposed into four components. The first term, $(1-\eta)\ln q_k$, captures the interaction between brokers and their analyst teams. If $\eta > 1$, for example, high-quality brokers

obtain lower marginal gains from additional analysts, implying substitutability between platform quality and analyst-team quality. The second term, $(\eta - \gamma) \ln \Omega_k$, reflects the interaction between analysts and their teams. As an example, $\eta > \gamma$ implies that the marginal contribution of additional analysts to broker composite quality is higher in high-quality teams. The third term, $(\gamma - 1) \ln a_{kx}$, captures the interaction among analysts of the same type within a team. When $\gamma < 1$, a given analyst type is more valuable when it is scarce within a team. Together with the second term, this implies complementarity across different types of analysts. Finally, the fourth term $\gamma \ln \omega_x$ indicates that, all else equal, high-quality analysts contribute more to brokers' composite quality.

Finally, we define an iterative algorithm to solve for the steady-state equilibrium given the parameter set:

Definition 2 (Solution Algorithm)**.** *Given the set of parameters* $\boldsymbol{\Theta}$*, the steady-state equilibrium can be solved in the following steps:*

1. ***Initial guess.*** *Guess an initial payoff matrix* $\boldsymbol{W}^{(0)}$*.*

2. ***Solve endogenous variables.*** *Solve analysts' money-metric expected value* $\widetilde{V}_{ix}\left(\boldsymbol{w}_x^{(0)}\right)$ *for all* $i \in \{0, 1, \ldots, N\}$ *and* $x \in \{1, \ldots, \overline{x}\}$ *by plugging* $\boldsymbol{W}^{(0)}$ *into (5). Plug the solution into (7) to solve analysts' job-switching fractions* $\mu_{ii'x}\left(\boldsymbol{w}_x^{(0)}\right)$ *for all* $i, i' \in \{0, 1, \ldots, N\}$ *and* $x \in \{1, \ldots, \overline{x}\}$*. Plug the fractions into (6), together with labor market clearing conditions (11) and (12), to solve the analyst allocation matrix* $\boldsymbol{A}\left(\boldsymbol{W}^{(0)}\right)$*. Plug the allocation matrix into (4) to compute the quality vector* $\boldsymbol{q}\left(\boldsymbol{W}^{(0)}\right)$*. Plug the quality vector into (2) to solve client shares* $\lambda_i\left(\boldsymbol{q}\left(\boldsymbol{W}^{(0)}\right)\right)$ *for all* $i \in \{1, \ldots, N\}$*. Finally, plug the solution into (3) to compute brokers' total surpluses* $S_i^T\left(\boldsymbol{q}\left(\boldsymbol{W}^{(0)}\right)\right)$ *for all* $i \in \{1, \ldots, N\}$*.*

3. ***Solve gradients.*** *Given the endogenous variables solved in step 2, solve the marginal effects of analysts' payoffs on their money-metric expected value,* $\frac{\partial \widetilde{V}_{kx}\left(\boldsymbol{w}_x^{(0)}\right)}{\partial w_{ix}^{(0)}}$ *for all* $i, k \in \{0, 1, \ldots, N\}$ *and* $x \in \{1, \ldots, \overline{x}\}$*, using (13) and (14). Solve the marginal effects of analysts on brokers' total surpluses,* $\frac{\partial S_i^T\left(\boldsymbol{q}\left(\boldsymbol{W}^{(0)}\right)\right)}{\partial a_{kx}\left(w_{kx}^{(0)}; \boldsymbol{w}_{\neg kx}^{(0)}\right)}$ *for all* $i, k \in \{1, \ldots, N\}$ *and* $x \in \{1, \ldots, \overline{x}\}$*, using (15), (16), (17), (18), and (19). Finally, solve the payoff sensitivities of brokers' marginal total surpluses from their own analysts,* $\frac{\partial}{\partial w_{ix}^{(0)}}\left(\frac{\partial S_i^T\left(\boldsymbol{q}\left(\boldsymbol{W}^{(0)}\right)\right)}{\partial a_{ix}\left(w_{ix}^{(0)}; \boldsymbol{w}_{\neg ix}^{(0)}\right)}\right)$ *for all* $i \in \{1, \ldots, N\}$ *and* $x \in \{1, \ldots, \overline{x}\}$*, using (20), (21), (22), and (23).*

4. ***Update payoffs.*** *Plug the endogenous variables solved in step 2 and the gradients solved in step 3 into (8), (9), and (10). Plug the Nash-bargaining surpluses obtained from (9) and (10), $S_{ix}^{A}\left(\boldsymbol{w}_{x}^{(0)}\right)$ and $S_{ix}^{B}\left(\boldsymbol{W}^{(0)}\right)$ for all $i \in \{1,\ldots,N\}$ and $x \in \{1,\ldots,\overline{x}\}$, further into (8). Use (8) to recompute each element of $\boldsymbol{W}^{(0)}$, except those in the exogenously given first row, and denote the resulting matrix as $\widetilde{\boldsymbol{W}}^{(1)}$. Finally, update the payoff matrix using a damping factor $\xi \in (0,1]$ such that $\boldsymbol{W}^{(1)} = (1-\xi)\,\boldsymbol{W}^{(0)} + \xi\widetilde{\boldsymbol{W}}^{(1)}$.*

5. ***Iteration.*** *Using the updated payoff matrix $\boldsymbol{W}^{(1)}$ as the new guess, repeat steps 2 through 4 to get $\boldsymbol{W}^{(2)}$, $\boldsymbol{W}^{(3)}$, and so on. The equilibrium payoff matrix $\boldsymbol{W}^{*}$ is reached when $\boldsymbol{W}^{(n+1)} = \boldsymbol{W}^{(n)}$.*

4 Quantitative Analysis

In this section, we take the model to the data and quantitatively examine how different labor market mechanisms affect brokers' composite quality and the resulting changes in stock price efficiency. Since the model characterizes a steady-state equilibrium, we focus on the cross-sectional allocation of analysts across brokers and use 2019 as the benchmark year for reporting the results.

4.1 Calibration and Estimation

Calibration We include all IBES brokers that cover at least one U.S. firm in the sample, resulting in 341 brokers in 2019. To prevent the job-switching matrix from becoming too sparse, we treat brokers with fewer than 15 analysts as part of the outside job option. This filtering leaves 73 large brokers, implying $N = 73$.[21] We classify analysts into three groups based on work experience: "High," "Middle," and "Low." Correspondingly, we set $\overline{x} = 3$ and calibrate the entries of the analyst endowment vector $\overline{\boldsymbol{a}}$ to match the shares of each analyst type in the overall analyst population.[22] Finally, we normalize the marginal

[21] Since the analyst labor market is highly concentrated, analysts employed by the excluded brokers account for only a small fraction of all analysts. Because our focus is on the effects of analyst labor market frictions on overall stock market efficiency, the quantitative impact of excluding these small brokers is negligible.

[22] We use work experience as the basis for our analyst classification for three reasons. First, analyst experience is cross-sectionally comparable across brokers because it depends only on the first year in which an analyst appears in the database and is unlikely to be confounded by platform quality. Second, as shown in our subsequent SMM estimation results, classifying analysts by work experience gener-

cost of brokerage service c to 1 and set the discount factor β to 0.9.

PPML estimation To obtain the inverse job-switching cost matrix $\widetilde{\boldsymbol{\tau}}$, we estimate the job-switching fraction equation (7) using the PPML regression with high-dimensional fixed effects (Correia et al., 2020):

$$\mu_{ii't} = \exp\left(\text{cons} + \delta_{it} + \delta_{i't} + \delta_{ii'p} + \varepsilon_{ii't}\right),$$

where $\mu_{ii't}$ is the fraction of analysts moving from employer i to employer i' at time t, including outside job options. Outside job options include unemployment, jobs in other industries, and jobs at excluded small brokers. The term $\exp\left(\delta_{it}\right)$ denotes origin-year fixed effects, which correspond to $\sum_{k=0}^{N}\exp\left(-\frac{\tau_{ik}}{\nu}\right)\widetilde{V}_k\left(\boldsymbol{w}\right)^{\frac{\beta}{\nu}}$ in the model. The term $\exp\left(\delta_{i't}\right)$ denotes destination-year fixed effects, which correspond to $\widetilde{V}_{i'}\left(\boldsymbol{w}\right)^{\frac{\beta}{\nu}}$ in the model. We divide the sample into two periods, pre- and post-2010, and use the origin-destination-period fixed effects $\exp\left(\delta_{ii'p}\right)$ as estimates of the entries $\exp\left(-\frac{\tau_{ii'}}{\nu}\right)$ in $\widetilde{\boldsymbol{\tau}}$. As in the model, we normalize the diagonal entries to one. To improve estimation precision, we include all sample years in the estimation and pool analysts across all types. The post-2010 estimates correspond to $\widetilde{\boldsymbol{\tau}}$ in 2019.

SMM estimation We estimate the remaining 86 parameters, collected in $\boldsymbol{\Theta}^{\mathbf{est}}$, using SMM. Specifically, we iteratively adjust the parameter values and run the solution algorithm in Definition 2 to match the model-implied analyst allocation matrix, $\boldsymbol{A}^{\mathbf{model}}\left(\boldsymbol{\Theta}^{\mathbf{est}}\right)$, with the allocation matrix observed in the data, $\boldsymbol{A}^{\mathbf{data}}$. The parameter estimates are given by:

$$\hat{\boldsymbol{\Theta}}^{\mathbf{est}} = \arg\min_{\boldsymbol{\Theta}^{\mathbf{est}}} \sum_{i=1}^{74}\sum_{x=1}^{3}\left[\ln\left(\frac{a_{ix}^{\text{model}}\left(\boldsymbol{\Theta}^{\mathbf{est}}\right)}{a_{ix}^{\text{data}}}\right)\right]^2,$$

where we use a log loss function to remove the influence of the moment scale. The estimation is overidentified, with 86 parameters used to match 222 moments.

Given that the model may admit multiple equilibria, we proceed in two steps. First, we construct well-motivated initial guesses based on the literature and fixed-effect regres-

ates clear quality differences across analyst types. Finally, work experience often shapes an analyst's role within a team, the bargaining process between the analyst and the broker, and the analyst's job mobility, all of which are closely aligned with the structure of the model.

sions. We then use the global optimization algorithm `patternsearch` to search for the optimum within a broad neighborhood around the initial guesses.

4.2 Estimation Results

Global parameters Table 1 reports the SMM estimates of the four global parameters. The parameter $\tilde{\theta}$, which governs the elasticity of client demand to markups, takes a relatively low value of 2.00. As shown in Proposition 1, this implies that brokers' composite quality plays a particularly important role in attracting clients, building market power, and extracting surplus. The estimates $\eta = 1.25$ and $\gamma = 0.80$ imply that platform quality and analyst-team quality are substitutes, while different types of analysts within a team are complements, as shown in equation (24). This matching mechanism is consistent with the evidence in sections 2.3 and 2.4, which documents that low-experience analysts realize larger gains in informativeness after moving to larger brokers and that the departure of high-experience analysts weakens the average market impact of their teams. Finally, the analyst preference dispersion parameter ν is estimated to be relatively high at 3.40, suggesting that non-pecuniary factors such as location, reputation, and specialization also play a non-negligible role in analysts' employer choices.

Table 1. SMM Estimates of Global Parameters

Parameter	Description	Value
$\tilde{\theta}$	Client demand elasticity	2.00
η	Broker-analyst interaction	1.25
γ	Within-team analyst interaction	0.80
ν	Analyst preference dispersion	3.40

Platform quality Figure 5 plots the model-implied composite quality q_i against the SMM-estimated platform quality b_i on a log scale, together with their corresponding kernel densities. Brokers in the estimation sample fall broadly into two groups. The vast majority of brokers are concentrated in the central cluster of the scatter plot, with relatively low and similar levels of platform quality. For these brokers, dispersion in composite quality is primarily driven by differences in analyst-team quality, as shown in equation (4).

By contrast, a small number of brokers in the upper-right corner of the scatter plot have exceptionally high platform quality and composite quality. Since the model implies substitutability between platform quality and analyst-team quality, these brokers rely less on high-quality analyst teams to attract clients. Instead, the approximately log-linear relationship suggests that dispersion in their composite quality is dominated by differences in platform quality. As shown by the kernel densities, these brokers form the long right tail of the quality distribution.

The long right tail of the quality distribution translates into high client concentration in the brokerage service market. Figure 6 plots the cumulative client share by ranking brokers according to their model-implied client shares λ_i. The top 10 brokers in the estimation sample account for roughly 80% of clients, while the remaining 20% are served by the other 63 brokers. This indicates that most clients are attracted by brokers' own quality rather than by high-quality analyst teams within brokers.

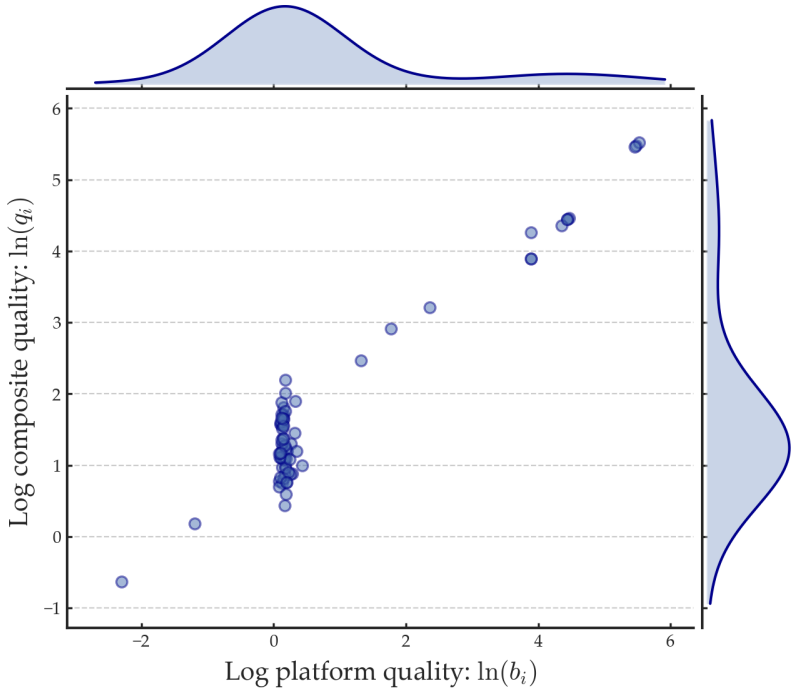


Figure 5. SMM Estimates of Platform Quality and Model-Implied Composite Quality.

Notes: This figure plots model-implied log composite quality, $\ln(q_i)$, against SMM-estimated log platform quality, $\ln(b_i)$. The right and top curves represent the corresponding kernel densities.

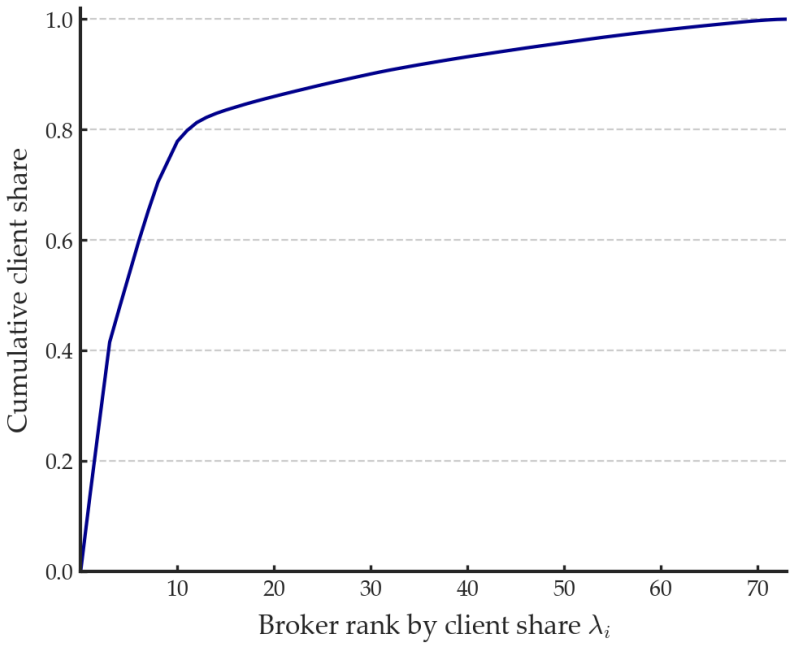


Figure 6. Model-Implied Client Share Concentration.

Notes: This figure plots cumulative client shares implied by the model, with brokers ranked in descending order by client share.

Analyst-type-specific parameters Table 2 reports the SMM estimates for the three sets of analyst-type-specific parameters. The analyst quality parameter ω_x increases strictly with analyst experience and shows clear differences across types. In particular, analysts in the high-experience group have a dominant quality level, more than an order of magnitude higher than those of the low- and middle-experience groups. The extremely high quality of high-experience analysts suggests that their expertise is more deeply embedded in the brokerage business and therefore less transferable to outside positions, resulting in a relatively low outside-option payoff $\overline{w}_x$. By contrast, analysts in the low- and middle-experience groups have higher outside option payoffs, possibly due to their greater career flexibility. Overall, the result is consistent with evidence in the labor economics literature that human capital is occupation-specific and accumulates through experience (e.g., Kambourov and Manovskii, 2009; Bagger et al., 2014).

For all analyst types, their relative bargaining power α_x is lower than that of brokers. Compared with analysts in the low- and middle-experience groups, high-experience analysts generate disproportionately more surplus for brokers. During the bargaining process, brokers therefore focus more on negotiations with high-experience analysts to ex-

tract more surplus net of wage payments, resulting in lower bargaining power for high-experience analysts than for low- and middle-experience analysts. This aligns with the survey evidence in Caldwell et al. (2026) that firms' propensity to bargain increases with job hierarchy.

Table 2. SMM Estimates of Analyst-Type-Specific Parameters

Parameter	Description	Value by Type		
		Low	Middle	High
ω_x	Analyst quality	17.225	50.453	693.101
$\overline{w}_x$	Outside job option payoff	0.009	0.009	0.001
α_x	Analyst bargaining power	0.250	0.250	0.125

To better understand the role of analyst-type-specific parameters in equilibrium, we first plot the model-implied job-switching fractions by analyst type in Figure 7. The lower quality and higher outside option payoffs of low- and middle-experience analysts result in job-switching rates between outside options and inside brokers at around 10%. In contrast, the corresponding fraction for more specialized high-experience analysts is only 2%. Among sample brokers, however, competition for analysts intensifies with analyst quality, leading to higher cross-broker job-switching rates for more experienced analysts. Yet, due to substantial job-switching costs in the financial analyst labor market, analyst mobility across sample brokers remains below 5% for all analyst types.

Figure 8 presents violin plots of brokers' marginal surplus from analysts, S^B_{ix}, and analysts' equilibrium payoff, w^*_{ix}, by analyst type. Despite having greater bargaining power than high-experience analysts, low- and middle-experience analysts receive equilibrium payoffs that are statistically lower than the marginal surplus retained by brokers. On the other hand, even with relatively lower bargaining power, high-experience analysts are still able to secure equilibrium payoffs that are statistically comparable to brokers' marginal surplus.

To better understand this pattern, it is worth noting that the surplus analysts obtain from bargaining is not their equilibrium payoff. Instead, it is the expected value that accounts for idiosyncratic preferences and potential future job switching, as given by S^A_{ix} in equation (9). The option value from idiosyncratic preference shocks and better future job opportunities thus accounts for a large share of the surplus of low- and middle-experience

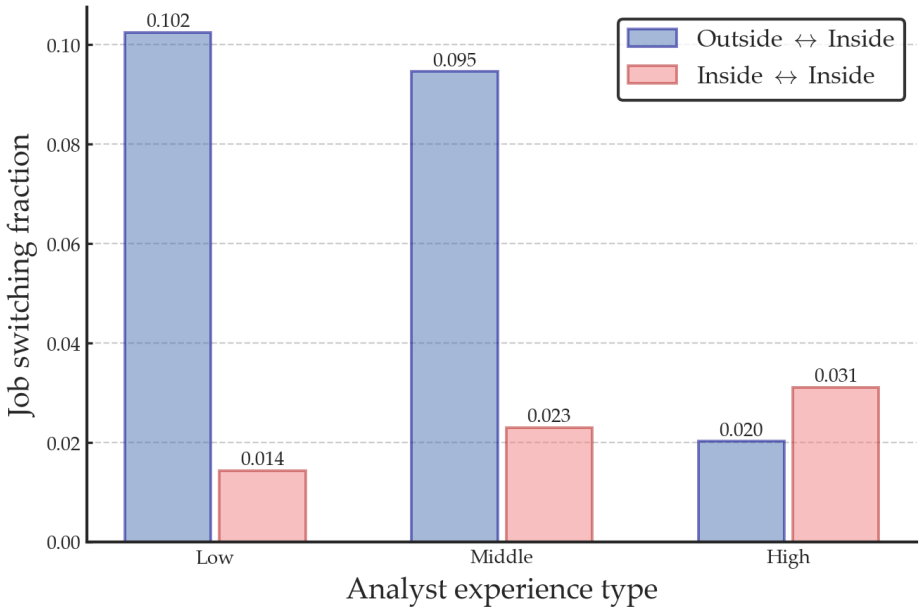


Figure 7. Model-Implied Job Switching Fractions by Analyst Type.

Notes: This figure shows model-implied job switching fractions by analyst type. For each analyst type, the blue bar reports the fraction of analysts moving between outside options and inside brokers. The red bar reports the fraction of analysts switching from one inside broker to another inside broker, excluding stayers.

analysts, whereas high-experience analysts place greater weight on the current wage flow.

Finally, consistent with differences in analyst quality, the equilibrium payoffs of high-experience analysts and the corresponding marginal surplus of brokers are both an order of magnitude higher than those for other analyst types. The much larger wage dispersion across analyst types than across brokers mirrors the literature on wage inequality (e.g., Abowd et al., 1999; Postel-Vinay and Robin, 2002; Bonhomme et al., 2019).

4.3 Model Fit

The quantitative fit of the model is summarized by Figure 9. Panel A presents a binned scatter plot of the overall model fit to the targeted moments. Except for a small number of analyst shares that are underestimated by the model, most points lie close to the 45-degree line. On average, the model fits the analyst allocation data well.

To examine the model's ability to match untargeted moments, we compare the model-implied total surplus attributable to analysts, S_i^T, with independently collected measures of brokerage revenue that are not used in the estimation. We manually collect publicly available financial filings for the brokerage houses in our sample, including SEC Form

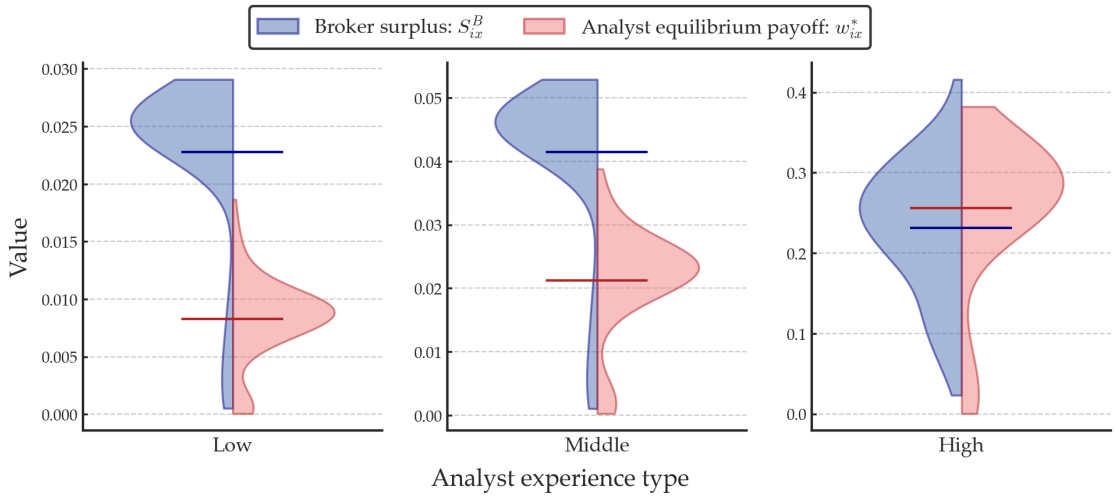


Figure 8. Model-Implied Broker Surplus and Equilibrium Payoffs by Analyst Type.

Notes: This figure presents violin plots of brokers' marginal surplus from analysts, S_{ix}^B, and analysts' equilibrium payoff, w_{ix}^*, by analyst type. Blue and red denote S_{ix}^B and w_{ix}^*, respectively. Curves and horizontal bars represent kernel densities and sample means, respectively.

10-K filings, annual reports, and other comparable regulatory or company disclosures. We then use ChatGPT 5.5 to extract revenue-side items from each filing, including total revenue, a narrow measure of brokerage revenue, and a broader measure of brokerage or capital-markets revenue when such items are separately reported and non-overlapping. Appendix C reports the prompt used for this extraction. To ensure data quality, all LLM-generated outputs are reviewed and verified by human researchers against the original filings.

Panel B shows a scatter plot of $\ln\left(S_i^T\right)$ against $\ln\left(\textit{BrokerageRevenue}\right)$, where *BrokerageRevenue* is defined as revenue from core brokerage services such as commissions. If S_i^T captures the total surplus attributable to analysts well, we expect a significantly positive relationship between $\ln\left(S_i^T\right)$ and $\ln\left(\textit{BrokerageRevenue}\right)$. The results in Panel B provide a consistent picture: $\ln\left(S_i^T\right)$ is significantly positively correlated with $\ln\left(\textit{BrokerageRevenue}\right)$ at the 5% level, with the coefficient estimate 0.346 implying that a 1% increase in revenue from core brokerage services is associated with a 0.346% increase in revenue attributable to analysts.

In Panel C, we replace *BrokerageRevenue* with *BroadBrokerageRevenue*, which includes core brokerage revenue as well as broader brokerage and capital-markets revenue items, such as underwriting, investment banking, and principal transactions, when they are sep-

arately reported. The scatter plot shows that $\ln(S_i^T)$ remains significantly positively correlated with $\ln(\textit{BroadBrokerageRevenue})$ at the 5% level. The coefficient estimate is smaller in magnitude: a 1% increase in the broader revenue measure is associated with a 0.255% increase in the model-implied surplus attributable to analysts. This economically weaker relationship is expected because the broader measure includes business lines that are less directly tied to analysts.

Finally, Panel D conducts a placebo test using the top-line total revenue of the brokerage firm, denoted by *TotalRevenue*. Unlike the narrow and broad brokerage revenue measures, *TotalRevenue* contains revenue from many other business lines that are unrelated or only weakly related to brokerage activity, such as retail banking and asset management. If S_i^T captures analyst-related surplus rather than the overall scale of the brokerage firm, it should not be systematically correlated with *TotalRevenue*. Consistent with this interpretation, the scatter plot shows a statistically insignificant relationship. This placebo result further validates the model's ability to capture the surplus created by analysts and the resulting endogenous allocation of analysts across brokers.

4.4 Composite Quality and Price Efficiency

In this section, we link the model-implied composite quality, q_i, to the empirical measure of broker-level aggregate informativeness of analyst reports. We expect higher broker composite quality to capture a stronger impact of its analysts on the stock prices of economically important firms.

Measuring informativeness Similar to the regression specification (1), we construct informativeness measures from market reactions to analysts' earnings forecast announcements. For each forecast announcement, we compute market-adjusted CARs over the $[-1, +1]$ window, where daily abnormal returns are stock returns minus the market return.

We first estimate informativeness at the analyst-broker-year level. For each analyst j, brokerage house i, and calendar year y, we use all forecasts issued by analyst j at brokerage i during years $y-1$ and y. The return-based informativeness measure is estimated separately for each analyst-broker-year using the following regression (Hilary and Hsu,

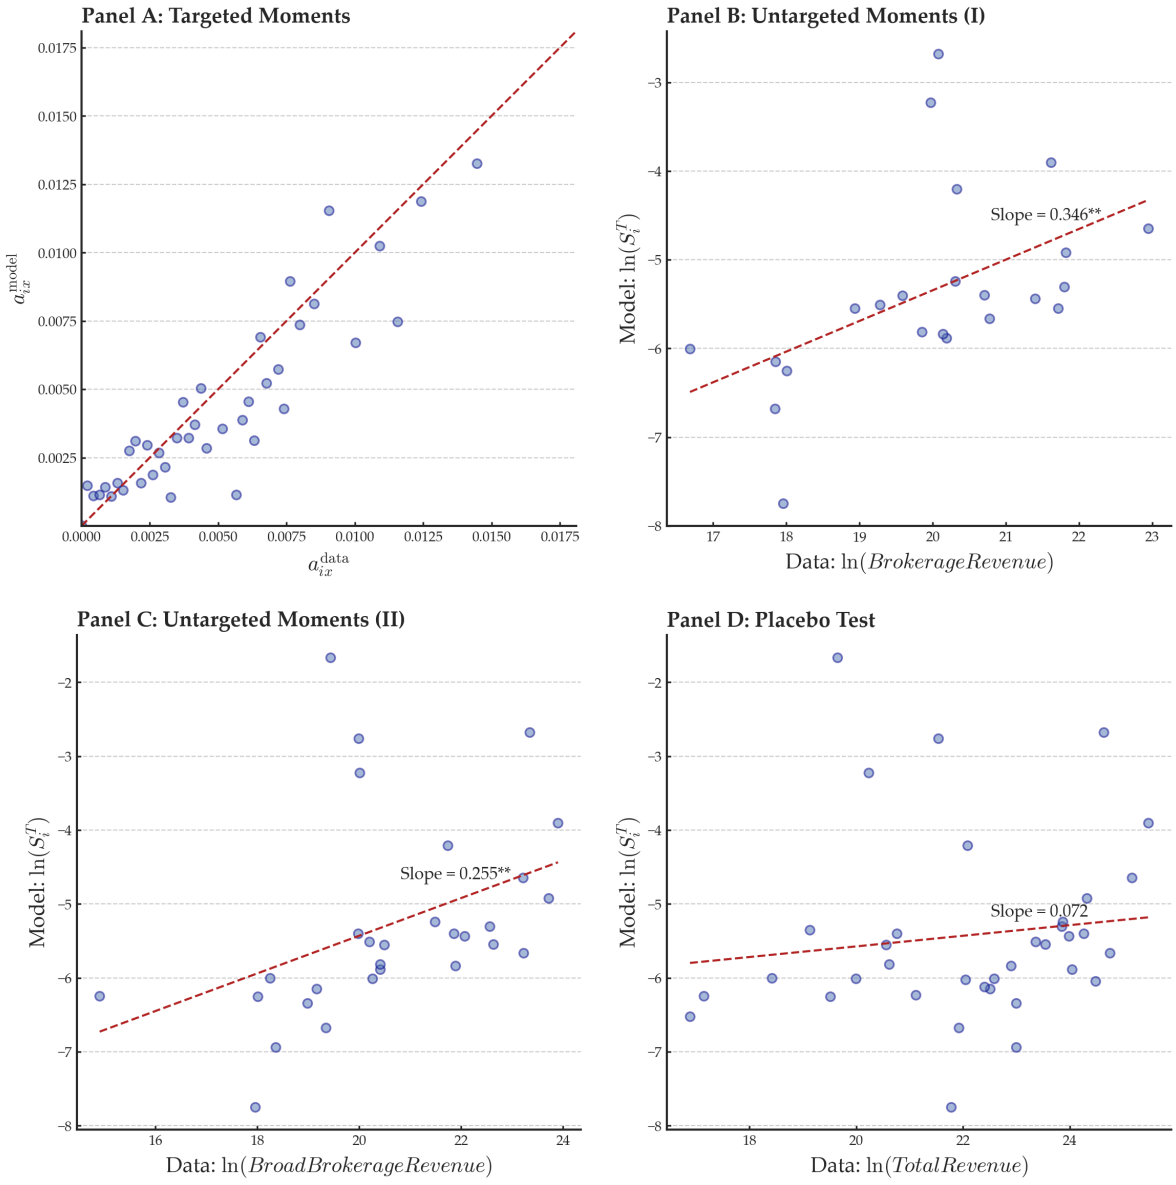


Figure 9. Model Fit.

Notes: This figure summarizes the model fit. Panel A presents a binned scatter plot of model-implied analyst allocation moments against their data counterparts, with the dashed red line indicating the 45-degree line. Panels B and C plot model-implied log total surplus attributable to analysts, $\ln\left(S_i^T\right)$, against log brokerage revenue and log broad brokerage revenue, respectively. Panel D plots $\ln\left(S_i^T\right)$ against log total revenue as a placebo test. Dashed red lines show fitted linear relationships.

2013):

$$CAR_{j,f,\tau} = \alpha^{CAR}_{i,j,y} + \beta^{CAR}_{i,j,y} \times SignedRevision_{j,f,\tau} + \varepsilon_{j,f,\tau},$$

where $CAR_{j,f,\tau}$ is the three-day market-adjusted cumulative abnormal return around analyst j's forecast announcement for firm f on date τ, and $SignedRevision_{j,f,\tau}$ is the change in analyst j's forecast relative to her prior forecast for the same earnings announcement, scaled by the stock price two days before the forecast date. A larger $\beta^{CAR}_{i,j,y}$ indicates that analyst j's forecast revisions elicit stronger price reactions when the analyst is affiliated with brokerage i.

As a robustness check, we also measure informativeness based on abnormal trading volume (Loh and Stulz, 2011). To construct this measure, we first compute abnormal trading volume using daily turnover. Daily turnover is defined as adjusted share volume divided by shares outstanding. For firm f on forecast date τ, abnormal volume is defined as

$$AVOL_{f,\tau} = \frac{\overline{Turnover}_{f,[-1,+1]}}{\overline{Turnover}_{f,[-61,-2]}} - 1,$$

where $\overline{Turnover}_{f,[-1,+1]}$ is firm f's average daily turnover over the three-day forecast-announcement window, and $\overline{Turnover}_{f,[-61,-2]}$ is firm f's average daily turnover over the pre-event benchmark window. The benchmark window begins 61 trading days before the forecast date and ends two trading days before the forecast date. We set the abnormal volume to missing if fewer than 30 valid observations are available in the benchmark window. We apply standard NASDAQ volume adjustments before computing turnover.[23]

The trading-volume-based informativeness measure is estimated as follows:

$$AVOL_{j,f,\tau} = \alpha^{AVOL}_{i,j,y} + \beta^{AVOL}_{i,j,y} \times \left|SignedRevision_{j,f,\tau}\right| + \varepsilon_{j,f,\tau}.$$

Because trading volume captures information arrival rather than the direction of news, we use the absolute value of signed forecast revisions in the volume regression. For both regressions, we require at least 36 valid forecast observations in the rolling two-year estimation window. All variables are winsorized at the 1st and 99th percentiles before estimation.

[23] Following prior work (e.g., Gao and Ritter, 2010), we divide NASDAQ share volume by 2.0 before February 1, 2001, by 1.8 during the remainder of 2001, and by 1.6 during 2002–2003. NASDAQ volume is left unadjusted from 2004 onward.

Finally, we aggregate analyst-level informativeness to the broker-year level by taking the weighted average of each analyst's $\beta_{i,j,y}^{CAR/AVOL}$:

$$BI_{i,y}^{CAR/AVOL} = \sum_{j=1}^{N_{i,y}} \omega_{i,j,y} \beta_{i,j,y}^{CAR/AVOL},$$

where the weight $\omega_{i,j,y}$ is based on the total market capitalization of the firms covered by analyst j. This weighting scheme gives greater importance to analysts whose coverage is economically more significant. Therefore, a higher value of $BI_{i,y}$ indicates that the information produced by brokerage i generates economically larger market reactions.

Validation To validate whether the model-implied composite quality also captures broker informativeness, we assume a proportional relationship and estimate the following specification:

$$BI_i^{CAR/AVOL} = b^{CAR/AVOL} \ln(q_i) + \varepsilon_i, \tag{25}$$

with the expectation that $b^{CAR/AVOL} > 0$. Figure 10 shows binned scatter plots of broker-level informativeness measures against log composite quality. In Panels A and B, we use return-based and trading-volume-based informativeness measures, respectively. Both measures are positively correlated with $\ln(q_i)$ at the 1% significance level. This indicates that composite quality in the model not only plays a crucial role in attracting clients and extracting surplus but also captures analysts' overall impact on the stock market within each broker. Put differently, high-quality brokers not only attract more clients but also produce analyst reports that generate stronger market reactions.

Economically, b^{CAR} = 1.399 implies that a 10% increase in composite quality raises the broker-level CAR sensitivity to forecast revisions by 0.1399. Evaluated at a one-standard-deviation increase in *SignedRevision* (0.0194), this corresponds to an incremental three-day CAR of about 0.27 percentage points. Applied to the average broker's covered market capitalization in 2019 ($9.2 trillion), this implies roughly $25 billion in additional market value affected.[24]

The trading-volume measure tells a similar story. With b^{AVOL} = 13.669, a 10% increase

[24] This calculation scales a one-standard-deviation revision across the average broker's entire coverage and is best read as an illustration of aggregate price discovery rather than a literal effect, since revisions occur firm by firm.

in composite quality raises the sensitivity of abnormal turnover to the absolute revision by 1.3669, so a one-standard-deviation revision implies an incremental abnormal-volume response of 0.027 (1.3669 $\times$ 0.0194). Because *AVOL* measures event-window turnover relative to its benchmark, this implies an extra 2.7 percentage points of abnormal turnover around the announcement relative to the average turnover in the past 60 days.

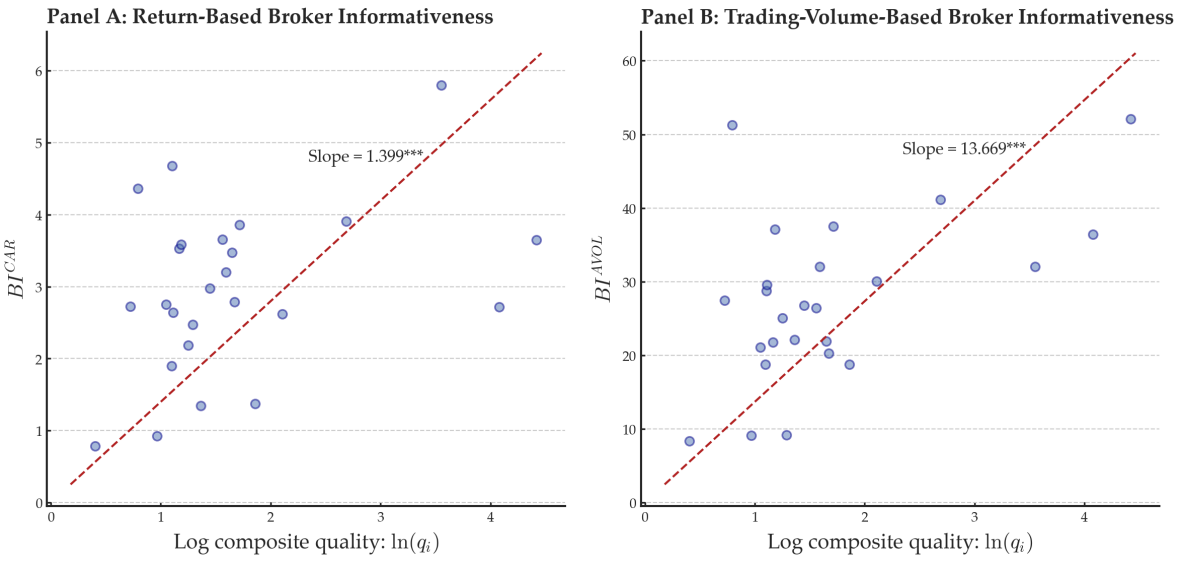


Figure 10. Composite Quality and Price Efficiency.

Notes: This figure presents a binned scatter plot of broker-level informativeness measures against model-implied log composite quality, $\ln(q_i)$. Panel A uses the return-based measure, BI^{CAR}, and Panel B uses the trading-volume-based measure, BI^{AVOL}. Dashed red lines show fitted proportional linear relationships.

4.5 Counterfactual Analysis

In this section, we quantify the role of different labor market mechanisms in shaping the aggregate impact of brokerage firms on price efficiency. Specifically, we examine how broker informativeness changes across counterfactual scenarios. Since BI_i captures the effect of a representative analyst at broker i on price efficiency, we focus on the analyst-share-weighted average of BI_i across all brokers. The proportional relationship in equation (25) implies that its percentage change in a counterfactual scenario relative to the baseline is given by:

$$\%\Delta\overline{BI}^{CAR/AVOL} = \frac{\overline{\ln(q)}^{\text{cf}} - \overline{\ln(q)}^{\text{base}}}{\overline{\ln(q)}^{\text{base}}}, \tag{26}$$

where $\overline{\ln(q)} = \sum_i \sum_x a_{ix} \ln(q_i)$ denotes the analyst-share-weighted average of model-implied log composite quality across brokers, and the superscripts cf and base refer to the counterfactual and baseline scenarios, respectively.

Price competition intensity We first quantify the effect of price competition intensity in the brokerage services on analyst labor market allocation and the resulting stock market efficiency by varying the client demand elasticity parameter $\tilde{\theta}$. Since the SMM estimate $\tilde{\theta} = 2.0$ implies low elasticity of client demand to markups, we increase $\tilde{\theta}$ to examine counterfactual scenarios in which clients are more sensitive to brokerage service prices. As shown in Figure 11, brokers' aggregate impact on price efficiency declines as price competition intensifies, with a 3.96% drop in the highly elastic case at $\tilde{\theta} = 10.0$.

According to Proposition 1, higher composite quality q_i and lower service markups $m(\lambda_i)$ are both crucial for brokers to attract clients and extract surplus. When $\tilde{\theta}$ is low, equation (18) implies that composite quality plays a relatively more important role in determining client shares λ_i. In this case, a high-quality analyst team is especially valuable because it enables a broker to secure greater market power, charge higher markups, and earn higher total surplus. As a result, brokers compete more intensively for analysts, and the quality matching between analysts and brokers improves. As $\tilde{\theta}$ increases, however, brokers shift their focus from competing in the analyst labor market to competing on service prices, which weakens the efficiency of quality matching and reduces the aggregate impact on price efficiency.

Job mobility Proposition 2 indicates that analyst job mobility, governed primarily by the idiosyncratic preference parameter ν and job-switching costs $\tau_{ii'}$, is a key determinant of analyst labor market allocation. In Panels A and B of Figure 12, we illustrate the counterfactual percentage changes in aggregate broker informativeness under different values of ν and $\tau_{ii'}$, respectively. The results suggest that stronger idiosyncratic preferences and lower job-switching costs play a non-negligible role in increasing brokers' aggregate impact on price efficiency. In particular, at $\nu = 10.0$, where idiosyncratic preferences play a dominant role in job choice, the impact on price efficiency increases by 5.10%. When we reduce job-switching costs to one tenth of their estimated level to approximate free mobility, the impact on price efficiency increases by 8.25%.

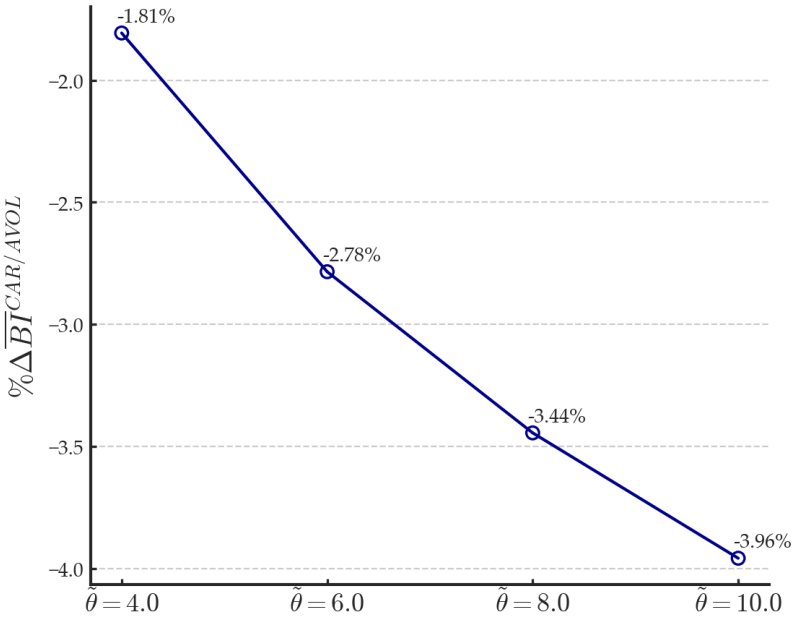


Figure 11. Counterfactual Price Competition Intensity.

Notes: This figure plots percentage changes in analyst-share-weighted average of broker informativeness, $\%\Delta\overline{BI}^{CAR/AVOL}$, under counterfactual changes in the client demand elasticity parameter $\tilde{\theta}$, where $\%\Delta\overline{BI}^{CAR/AVOL}$ is given by equation (26).

To understand the mechanism, recall that the SMM estimates imply substitutability between platform quality and analyst-team quality. As a result, analysts at brokers with dominant platform quality make a lower marginal contribution to composite quality. By contrast, mid- and lower-tier brokers rely more heavily on high-quality analyst teams. Under weak idiosyncratic preferences and high job-switching costs, however, analyst inflows into mid- and lower-tier brokers are impeded. Because analysts from different experience groups are complementary within teams, low job mobility translates into less effective team composition, thereby dampening brokers' aggregate impact on price efficiency.

As shown in equation (23), increasing the idiosyncratic preference parameter ν uniformly lowers the elasticity of job-switching fractions with respect to analyst payoffs. As shown in equations (5) and (9), reducing job-switching costs $\tau_{ii'}$ increases analysts' option value from future job switching and lowers their current surplus. Both mechanisms facilitate the inflow of complementary analysts into mid- and lower-tier brokers.

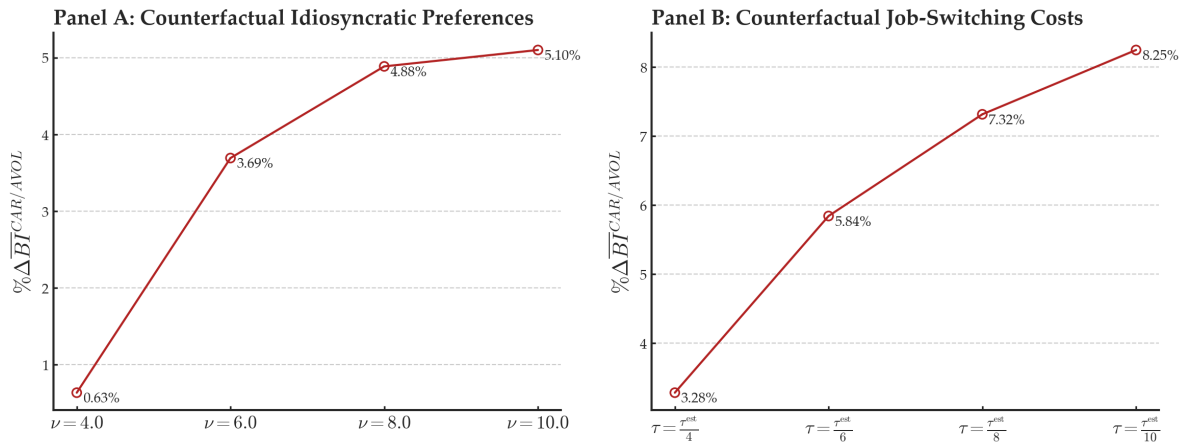


Figure 12. Counterfactual Job Mobility.

Notes: Panels A and B of this figure illustrate the counterfactual percentage changes in aggregate broker informativeness under different values of ν and $\tau_{ii'}$, respectively, where $\%\Delta\overline{BI}^{CAR/AVOL}$ is given by equation (26).

Bargaining power Equation (8) in Proposition 3 shows that analysts' relative bargaining power α_x affects their labor market allocation by directly affecting wage determination. Our SMM estimates suggest that analysts are in a disadvantaged position when bargaining with brokers. To quantify the effect of this asymmetric bargaining power on aggregate broker informativeness, we examine counterfactual scenarios in which the two parties have equal bargaining power. As shown in Figure 13, when we increase only the bargaining power of high-experience analysts to 0.5, brokers' aggregate impact on price efficiency increases by 3.10%. A larger increase of 3.46% is obtained when we set the bargaining power of all analyst types uniformly to 0.5. In contrast, increasing only the bargaining power of low- and middle-experience analysts reduces the aggregate impact by 0.16%.

To understand the differences, recall that high-experience analysts have a distinctively higher quality than other analyst types. As the bargaining power of high-experience analysts increases, a larger share of the broker surplus they create is translated into their equilibrium payoffs. This stronger link between analysts' payoffs and their marginal contributions to brokers facilitates the reallocation of high-experience analysts toward brokers with greater demand for high-quality analyst teams, thereby raising brokers' aggregate impact on price efficiency.

For low- and middle-experience analysts, whose quality is much lower, increasing

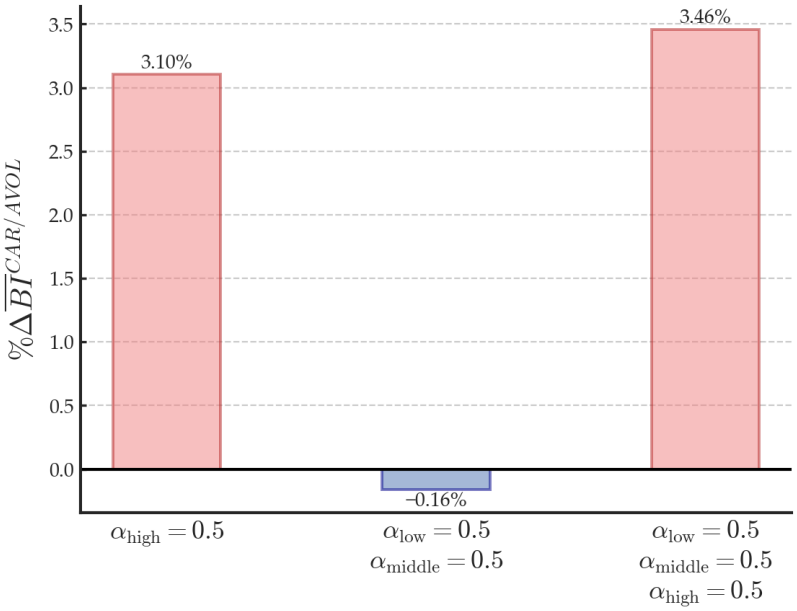


Figure 13. Counterfactual Bargaining Power.

Notes: This figure plots percentage changes in analyst-share-weighted average of broker informativeness, $\%\Delta\overline{BI}^{CAR/AVOL}$, under counterfactual changes in analyst bargaining power $\boldsymbol{\alpha}$, where $\%\Delta\overline{BI}^{CAR/AVOL}$ is given by equation (26).

their bargaining power has only a limited effect. If anything, the larger gap in bargaining power relative to high-experience analysts makes it harder for brokers to form balanced analyst teams, which slightly reduces the aggregate price impact. When the bargaining power of high-experience analysts is also increased, however, stronger team complementarity generates an additional gain in price efficiency.

Matching mechanisms Taken together, the above counterfactual analyses suggest that adjustments to analyst labor-market frictions have only a modest effect on aggregate broker informativeness, with changes of less than 10%. To understand the reason, recall that Figure 5 shows that a small number of brokers have exceptionally high platform quality b_i. Because platform quality and analyst-team quality are substitutes, the marginal effect of gaining or losing any analyst is limited for these brokers. As a result, the counterfactual effects of changing analyst labor market frictions are primarily driven by changes in analyst-team quality within mid- and lower-tier brokers, whose price impact is more sensitive to team composition. However, because the market for brokerage services is highly concentrated, as shown in Figure 6, these brokers contribute relatively little to aggregate

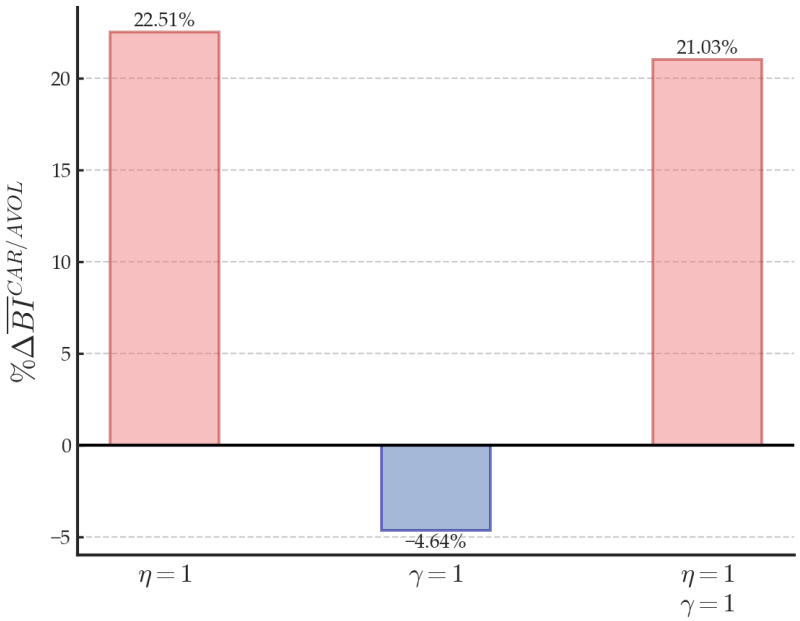


Figure 14. Counterfactual Quality Matching.

Notes: This figure plots percentage changes in analyst-share-weighted average of broker informativeness, $\%\Delta\overline{BI}^{CAR/AVOL}$, under counterfactual changes in analyst-broker matching mechanisms, where $\%\Delta\overline{BI}^{CAR/AVOL}$ is given by equation (26).

price efficiency.

As a final set of counterfactual analyses, we remove the quality-matching mechanisms between brokers and analysts to quantify their impact on price efficiency. In the first counterfactual, we remove the substitutability between platform quality and analyst-team quality by setting $\eta = 1$. In this case, clients evaluate analyst quality independently of platform quality, so that the contribution of analyst teams at top-tier brokers is no longer discounted by the brokers' own quality. Since top-tier brokers are the main contributors to aggregate brokerage informativeness, the aggregate impact increases substantially, by 22.51%. In comparison, removing complementarity among different analyst types within teams by setting $\gamma = 1$ reduces brokers' aggregate impact on price efficiency. Again, because team synergies play a more important role in mid- and lower-tier brokers, the magnitude of this effect is much smaller, at 4.64%. Finally, simultaneously shutting down the two quality-matching mechanisms raises aggregate broker informativeness by 21.03%.

4.6 Policy Implications

Our estimates speak to an active debate about how to strengthen the informativeness of sell-side research. The central lesson is that price efficiency depends on the labor-market arrangements through which it is produced. We organize the discussion around four policy margins, beginning with the one tied to our largest counterfactual effect.

Improving transparency in research pricing. Our largest counterfactual effect concerns the substitutability between platform quality and analyst-team quality. Brokers with strong platform quality attract clients even when the marginal contribution of their analyst teams is low. Their client base may rest on reputation, corporate access, investment-banking relationships, or bundled services, so an additional high-quality analyst has only a small effect on composite quality. This weakens the link between client demand and analyst-team formation. If clients cannot distinguish a strong analyst team from a strong platform, large platforms may underinvest in analyst talent. When this substitutability is removed, the success of large brokers depends more on analyst-team quality. As a result, they compete more intensively in the analyst labor market, allowing talent to flow more easily to where it is most valuable.

These results have direct implications for the regulation of research payments. The key issue is whether clients can observe and pay for the value of analyst research separately from the value of the brokerage platform that delivers it. Historically, sell-side research has often been bundled with broader brokerage and investment-banking services, so research departments were pressured to justify themselves through trading commissions and banking revenues rather than through the quality of their output (Bradshaw et al., 2017). This form of bundling remains especially relevant in the U.S., where research payments have not been broadly unbundled. Bundled "soft-dollar" payments remain permissible under the Section 28(e) safe harbor of the Securities Exchange Act of 1934, subject to best-execution and disclosure obligations.[25] Although temporary SEC relief allowed U.S. broker-dealers to accept the separate research payments required by MiFID II for European managers, that relief expired in 2023 (SEC, 2017, 2019).

[25] When research is funded opaquely through bundled commissions or embedded in broader brokerage relationships, incentives can drift toward trade generation, relationship maintenance, and other platform rents rather than toward pure research quality.

This institutional setting is precisely where the model's substitution channel is most likely to matter. When research and execution are bundled, clients may direct business to dominant platforms for reasons unrelated to research informativeness. As a result, the marginal return to competing through analyst quality is lower for the brokers with stronger platforms. In the debate over the expiration of the SEC relief, both industry participants and Commissioner Uyeda warned that restricting bundled payments could reduce the supply of research and weaken price discovery, especially for smaller issuers (SIFMA, 2023; Uyeda, 2023). Permissive bundling can preserve the platform-analyst substitution that our counterfactuals identify as the largest drag on price efficiency. It can also disadvantage mid- and lower-tier brokers, where analyst talent has the highest marginal value but platform rents are weaker.

Europe provides the opposite experiment. MiFID II, implemented in 2018, required asset managers to separate research payments from execution payments and to justify research spending more explicitly. Guo and Mota (2021) show that this reform reduced the quantity of analyst coverage but improved its quality. Inaccurate analysts exited, while remaining analysts produced more accurate research as asset managers became more selective and competition shifted toward research quality.[26] This pattern is consistent with the model. By making research payments more transparent, unbundling weakens the ability of strong platforms to monetize bundled relationships and increases the reward to analyst-team quality.

Our results therefore suggest caution about the recent move back toward bundling. The UK has reintroduced a joint-payment option for research and execution, first for MiFID investment firms (FCA, 2024) and then for UK fund managers (FCA, 2025). The EU Listing Act has moved in a similar direction (EU, 2024). These reforms are motivated by a legitimate goal: reviving research coverage, especially for smaller listed firms (Kent, 2023). Importantly, both jurisdictions retained transparency requirements. Firms using bundled payments must disclose their payment approach, set and disclose research budgets, manage conflicts of interest, and periodically assess the quality and value of the research they purchase (FCA, 2025; EU, 2024). These guardrails are central from the perspective of the model, because they help ensure that research is priced according to infor-

[26] Notably, Guo and Mota (2021) find that the coverage decline was concentrated among large firms rather than the smaller companies that unbundling was feared to harm. This is consistent with regulatory evidence that SMEs were not disproportionately affected (FCA, 2019; Amzallag et al., 2021).

mativeness rather than merely bundled with platform access.

The policy implication is therefore not that one payment mechanism is always optimal. Rather, the key safeguard is transparency. Whether regulators maintain permissive bundling, as in the U.S., or relax strict unbundling, as in the EU and UK, research payments should remain closely tied to the informativeness of research. If re-bundling weakens that link, it risks restoring platform rents and disproportionately benefiting larger brokers. If transparency is preserved, bundled payments are less likely to undermine competition in analyst quality and less likely to divert talent away from the mid- and lower-tier brokers whose competitiveness depends most on strong analyst teams.

Reducing frictions to analyst mobility. Cutting job-switching costs raises analysts' aggregate price impact by promoting the flow of analysts to mid- and lower-tier brokers where their marginal informational contribution is greatest. In addition, because analysts are complements within teams, mobility barriers also impede the formation of productive teams by limiting the grouping of complementary analysts.

A broader labor literature speaks to this margin. Enforceable non-compete agreements reduce mobility and are associated with lower wages where enforcement is stronger (Starr et al., 2021), especially for workers with firm-specific skills and narrow specializations in knowledge-intensive occupations (Marx et al., 2009; Jeffers, 2024). These features closely describe senior, specialized analysts entrenched at incumbent platforms. The relevant policy margin is calibration rather than elimination. Firms have legitimate interests in protecting investments in client relationships and proprietary information, but survey evidence shows many non-competes are imposed indiscriminately, including workers with no such interest at stake and often only after a job has been accepted (Starr et al., 2021). This points to narrower instruments, such as capping the duration of non-compete agreements and garden-leave periods, limiting enforceability to higher-paid employees, and requiring compensation during restricted periods. Because these arrangements are governed largely by state law in the U.S., enforceability varies widely, leaving ample room for reform. The aim is not turnover for its own sake, but to allow talent to reach the organizations where its marginal value is highest.

Strengthening competition in research provision. The gains from greater mobility are concentrated among small- and medium-sized research providers. Large platforms attract clients through advantages less directly tied to analyst teams, such as established networks, superior information access, and reputational capital, while mid- and lower-tier brokers must compete more heavily on team quality. This reflects the logic of two-sided markets, in which a dominant intermediary that controls access to one side can extract rents from producers on the other (Rochet and Tirole, 2003). In the brokerage market, these rents are reinforced by concentration. Such concentration can degrade research quality by weakening competition among analysts (Hong and Kacperczyk, 2010).

Two interventions follow. The first is to lower the fixed costs of producing and distributing research, which fall disproportionately on smaller brokers. The central research platform proposed by the UK Investment Research Review is one such measure, sourcing and freely disseminating research on smaller listed firms (Kent, 2023). By socializing the distribution infrastructure that large brokers supply internally, such a platform would narrow the non-research advantages that sustain the substitution channel.

The second intervention is to widen retail access to research. Independent, retail-facing research has become a competitive force in its own right, and analysts respond to its preemption with bolder, more disaggregated forecasts (Drake et al., 2023). This disciplines established brokers by raising the cost of producing uninformative research. Broadening its reach reduces smaller providers' dependence on dominant platforms for distribution, raising the return to independent research.

Increasing pay transparency and output visibility. Equalizing high-experience analysts' bargaining power with brokers' raises aggregate informativeness. Although these analysts are highly productive, they bargain from a relatively weak position. Policies that improve their knowledge of outside opportunities could therefore align compensation more closely with marginal contribution and facilitate their movement toward brokers where their contribution to research quality is greatest.

Greater external wage transparency provides one such policy lever. Workers often underestimate wages available at other employers and consequently undervalue their outside options (Jäger et al., 2024). Requiring credible salary ranges in job advertisements and publishing comparable compensation benchmarks by experience and performance

would reduce the informational advantage of incumbent employers. Consistent with this mechanism, cross-firm pay transparency encourages workers to pursue better-paying opportunities, negotiate higher compensation, and intensifies competition among employers (Cullen, 2024). In our framework, such policies increase the perceived value of analysts' alternative job choices, thereby directing analysts toward the teams where their marginal contribution is greatest.

For specialized workers such as financial analysts, however, wage transparency alone may be insufficient: Prospective employers must also be able to assess individual productivity. Complementary policies should therefore make analyst output verifiable and portable across employers. Standardized records of forecast accuracy, market responses to research reports, investor votes, and other auditable indicators of research quality would help competing brokers separate an analyst's contribution from the reputation and resources of her current platform. Public analyst rankings already serve this function by providing visible signals of ability to clients and prospective employers (Groysberg et al., 2011). Making such performance measures more systematic and portable would encourage brokers to compete for productive analysts and strengthen analysts' bargaining power.

5 Conclusion

Sell-side analysts are central information producers whose research contributes directly to price efficiency, yet the labor market that allocates them across brokerage platforms has received little attention. This paper treats this allocation as part of the information-production technology of capital markets and examines how different labor-market mechanisms shape it.

Three facts motivate the analysis. First, analyst-broker matching matters. Using brokerage M&As, we find that junior analysts experience larger gains in forecast informativeness after moving to larger platforms, pointing to substitutability between platform quality and analyst ability. Second, analyst-team composition matters. Using senior analyst retirements, we find that the informativeness of remaining junior analysts declines after a senior colleague leaves, pointing to complementarity among different types of analysts within teams. Finally, analyst job mobility has declined sharply over time, es-

pecially among highly experienced analysts who often play a critical role within teams, suggesting that barriers to labor-market matching have become more severe.

To quantify how these forces affect aggregate price efficiency in equilibrium, we develop and estimate a structural model in which oligopolistic brokers compete for analysts to attract client demand, forward-looking analysts choose employers subject to mobility costs and idiosyncratic preferences, and wages are determined through Nash-in-Nash bargaining. The estimated model reproduces the targeted allocation of analysts, and its implied broker surplus and composite quality align with untargeted measures of brokerage revenue and analyst forecast informativeness.

Our central insight is that labor-market frictions limit price efficiency mainly by impeding the flow of analysts toward the mid- and lower-tier brokers that rely more heavily on high-quality analyst teams, and within-team complementarity further compounds the shortfall. Any counterfactual that facilitates this flow thus improves the efficiency of analyst allocation, and vice versa. For example, reducing job-switching costs to one tenth of their estimated level raises analysts' aggregate contribution to price efficiency by 8.25%, equalizing analysts' bargaining power with brokers' increases it by 3.46%, while intensifying price competition in the brokerage-service market lowers it. The matching technology itself is first-order: Removing the platform-team substitution raises aggregate informativeness by 22.51%, our largest effect.

These results have clear policy implications. Because the substitution between platform quality and analyst-team quality is the dominant source of lost price efficiency, policy should strengthen transparency in research pricing. Clients should be able to evaluate and pay for the informativeness of analyst research separately from the platform through which that research is delivered. The remaining policy margins reinforce the same goal of directing talent to where its marginal contribution is greatest: better-calibrated mobility restrictions, lower fixed costs and wider retail access for smaller research providers, and greater external wage transparency paired with portable performance records that make analyst value visible to rival employers.

Several limitations point to promising avenues for future research. First, our analysis characterizes a steady-state equilibrium and therefore focuses on the cross-sectional allocation of analysts. Future work could model the transition dynamics underlying the decline in labor-market mobility and examine how shocks propagate as the market ad-

justs. Second, we treat broker platform quality and analyst ability as exogenous endowments, abstracting from brokers' investments in research infrastructure and analysts' accumulation of human capital. Endogenizing these choices would connect labor-market design to the long-run formation of the quality differences that our model takes as given. Finally, our framework extends beyond financial analysts. Similar allocation problems arise whenever workers produce intangible outputs that support demand in profitable product markets, suggesting that labor-market design may be an important determinant of market efficiency more broadly.

Appendix

A Additional Tables and Figures

Table A1. Brokerage Mergers and Acquisitions

Date	Type	Target	Acquirer	N_T	N_A	N_M
31/05/1997	Merger	Dean Witter Discover	Morgan Stanley	39	80	6
09/01/1998	Acquisition	Principal Financial Securities	EVEREN Capital	19	33	2
17/02/1998	Acquisition	Jensen Securities	D.A. Davidson	5	9	4
03/11/2000	Acquisition	Donaldson Lufkin & Jenrette	Credit Suisse First Boston	79	156	23
31/12/2000	Merger	Chase Manhattan	J.P. Morgan	50	72	33
31/05/2001	Acquisition	ING Financial Markets	ABN Amro	53	91	20
27/07/2001	Acquisition	Robinson Humphrey	SunTrust Robinson-Humphrey	26	32	3
01/07/2005	Merger	Parker/Hunter	Janney Montgomery Scott	4	14	3
28/02/2007	Acquisition	Ryan Beck	Stifel Financial	17	55	6
31/08/2007	Acquisition	Cochran Caronia Securities	Fox-Pitt Kelton	3	24	3
30/05/2008	Acquisition	Bear Stearns	JPMorgan Chase	84	115	19
15/02/2009	Acquisition	Pacific Growth Equities	Wedbush Morgan Securities	16	18	8
01/09/2009	Acquisition	Tristone Capital	Macquarie	7	26	4
01/07/2010	Acquisition	Thomas Weisel Partners	Stifel Financial	34	64	21
02/04/2012	Acquisition	Morgan Keegan	Raymond James	24	87	5
17/12/2012	Acquisition	Caris	B. Riley	13	16	5
03/09/2014	Acquisition	Pacific Crest Securities	KeyCorp	19	28	13
29/02/2016	Acquisition	Simmons	Piper Jaffray	13	44	9
30/04/2019	Acquisition	FIG Partners	Janney Montgomery Scott	6	11	4
06/01/2020	Merger	Sandler O'Neill	Piper Sandler	21	44	17

Notes: This table lists the brokerage M&A events used in the identification strategy in section 2.3. We start from the M&A events used in prior studies (Hong and Kacperczyk, 2010; Kelly and Ljungqvist, 2012) and supplement them with hand-collected events from public sources. For acquisitions, we require the acquiring broker to be larger than the target broker before the transaction. For mergers, we define the moving direction as from the smaller broker to the larger broker. Broker size is measured by the number of distinct analysts issuing IBES forecasts during the 12 months before the event date. The pre-event window is the 12 months ending one month before the event date, and the post-event window is the 12 months beginning one month after the event date. The one-month gap allows for differences between announcement dates, completion dates, and the operational integration of research teams. We require each event to involve at least one moving analyst, defined as an analyst who issues forecasts at the target/smaller broker before the event and at the acquiring/larger broker after the event. N_T and N_A denote the number of analysts at the target/smaller broker and acquiring/larger broker, respectively, during the pre-event window. N_M denotes the number of analysts who move from the target/smaller broker to the acquiring/larger broker after the event.

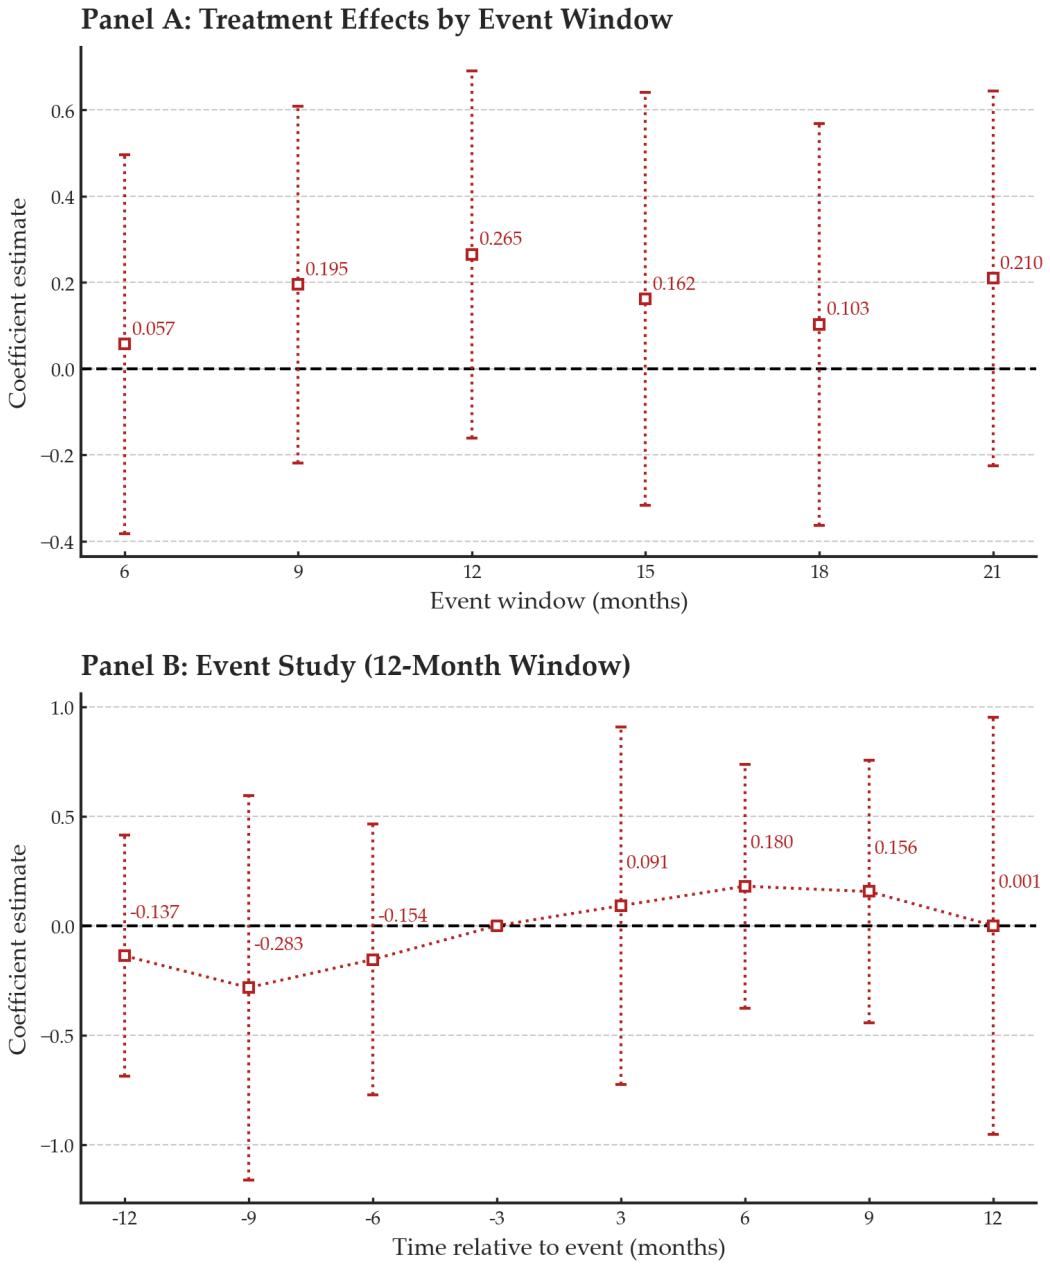


Figure A1. Brokerage M&As and Analyst Forecast Informativeness: Placebo Tests.

Notes: This figure reports placebo estimates obtained by assigning each brokerage merger a pseudo-event date 24 months before the actual merger date. The placebo sample includes analysts who remain at the target brokerage and issue forecasts on both sides of the pseudo-event date. Unlike the main analysis, analysts are not required to subsequently move to the acquiring brokerage. Treated analysts are those with below-median tenure during the placebo pre-period, while analysts with above-median tenure form the control group. Panel A plots the coefficient on *SignedRevision* × *Treatment* × *Post* across alternative symmetric event windows, excluding one month on either side of the pseudo-event date. Panel B reports the corresponding quarterly event-study coefficients for the 12-month window, with the quarter immediately preceding the pseudo event omitted as the reference point. All regressions include analyst-event-firm and year-month fixed effects. Standard errors are clustered by merger event. Point estimates are labeled, and vertical bars represent 95% confidence intervals.

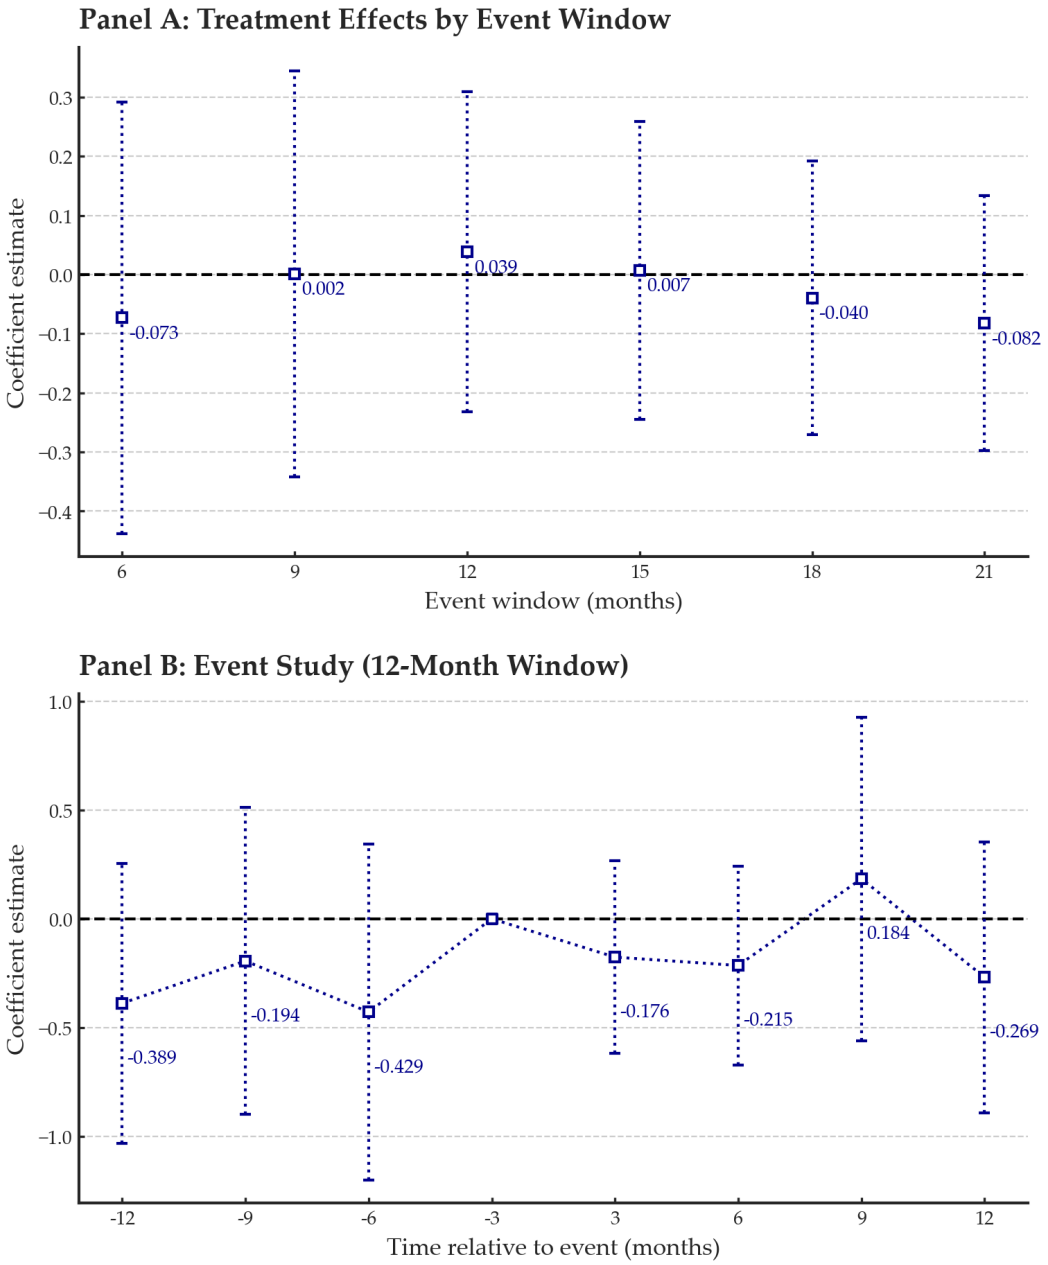


Figure A2. Senior Colleague Retirements and Analyst Forecast Informativeness: Placebo Tests.

Notes. This figure reports placebo tests for the effect of senior analyst retirements on the informativeness of remaining junior analysts' forecast revisions. Placebo events are senior analysts who satisfy the same eligibility thresholds as the main analysis (older than 55 and at least 20 years of financial-industry experience) but do not retire. The pseudo-event date is set to June 1 of the first eligible year. Using the same stacked DiD design, we estimate *SignedRevision* × *Treatment* × *Post*, where treatment refers to the remaining junior analysts (tenure below 5 years) covering the same Fama-French 30 industry as the pseudo retiree. All specifications include analyst-event-firm and year-month fixed effects, with standard errors clustered at the event level. Panel A plots the triple interaction terms from six separate regressions, each using a symmetric pre/post estimation window of the length shown on the horizontal axis. Panel B reports event-study estimates for the 12-month window, relative to the quarter immediately before the event. Point estimates are labeled, and vertical bars represent 95% confidence intervals.

B Proofs

B.1 Proof of Proposition 1

By the law of large numbers, the fraction of clients choosing broker i equals the corresponding choice probability:

$$\begin{aligned}
\lambda_i &= \Pr\left(\arg\max_{i'}\left\{f_{i'}^{-\kappa}\varepsilon_{i'}(o)\right\} = i\right) \\
&= \int_0^\infty \prod_{i'\neq i} \Pr\left(f_{i'}^{-\kappa}\varepsilon_{i'} \leq u\right) \mathrm{dPr}\left(f_i^{-\kappa}\varepsilon_i \leq u\right) \\
&= \int_0^\infty \prod_{i'\neq i} G_{i'}\left(uf_{i'}^{\kappa}\right) \mathrm{d}G_i\left(uf_i^{\kappa}\right) \\
&= \int_0^\infty \theta u^{-\theta-1} q_i f_i^{-\tilde{\theta}} \exp\left(-u^{-\theta}\sum_{i'=1}^{N} q_{i'} f_{i'}^{-\tilde{\theta}}\right) \mathrm{d}u \\
&= \frac{q_i f_i^{-\tilde{\theta}}}{\sum_{i'=1}^{N} q_{i'} f_{i'}^{-\tilde{\theta}}}.
\end{aligned} \tag{27}$$

The first-order condition (FOC) for broker i is given by:

$$\frac{\partial \lambda_i}{\partial f_i}(f_i - c) + \lambda_i = 0.$$

Using equation (27), we have:

$$\frac{\partial \lambda_i}{\partial f_i} = -\frac{\tilde{\theta}\lambda_i(1-\lambda_i)}{f_i}.$$

Substituting into the FOC yields:

$$f_i = \frac{\tilde{\theta}(1-\lambda_i)}{\tilde{\theta}(1-\lambda_i) - 1}c = \frac{\sigma(\lambda_i)}{\sigma(\lambda_i) - 1}c = m(\lambda_i)\, c.$$

Plugging f_i into λ_i, we obtain:

$$\lambda_i = \frac{q_i m(\lambda_i)^{-\tilde{\theta}}}{\sum_{i'=1}^{N} q_{i'} m(\lambda_{i'})^{-\tilde{\theta}}}, \quad \forall\, i \in \{1, \ldots, N\},$$

which corresponds to equation (2).

Finally, the total surplus is given by:

$$S_i^T(\boldsymbol{q}) = \lambda_i(\boldsymbol{q})\left[m\left(\lambda_i(\boldsymbol{q})\right) - 1\right]c = \frac{\lambda_i(\boldsymbol{q})}{\sigma\left(\lambda_i(\boldsymbol{q})\right) - 1}c = \frac{\lambda_i(\boldsymbol{q})}{\tilde{\theta}\left[1 - \lambda_i(\boldsymbol{q})\right] - 1}c,$$

which is equation (3). □

B.2 Proof of Proposition 2

The Gumbel distribution assumption implies that:

$$\begin{aligned}
\mathbb{E}\left[\max_{i'}\left\{\beta V_{i'xt+1} - \tau_{ii'} + \nu\varphi_{i'xt}\right\}\right] &= \int_{-\infty}^{\infty} u \,\mathrm{d}\prod_{i'=0}^{N}\Pr\left(\beta V_{i'xt+1} - \tau_{ii'} + \nu\varphi_{i'xt} \leq u\right) \\
&= \int_{-\infty}^{\infty} u \,\mathrm{d}\prod_{i'=0}^{N} F\left(\frac{u - \beta V_{i'xt+1} + \tau_{ii'}}{\nu}\right) \\
&= \frac{\Phi_{ix}}{\nu}\int_{-\infty}^{\infty} u \exp\left(-\frac{u}{\nu} - \overline{\gamma}\right)\exp\left(-\Phi_{ix}\exp\left(-\frac{u}{\nu} - \overline{\gamma}\right)\right)\mathrm{d}u,
\end{aligned}$$

with $\Phi_{ix} = \sum_{i'=0}^{N}\exp\left(\frac{\beta V_{i'xt+1} - \tau_{ii'}}{\nu}\right)$.

Using the change of variables $\exp\left(-\frac{u}{\nu} - \overline{\gamma}\right) = y$, we have $u = -\nu\left(\ln y + \overline{\gamma}\right)$ and thus:

$$\frac{\Phi_{ix}}{\nu}\int_{-\infty}^{\infty} u \exp\left(-\frac{u}{\nu} - \overline{\gamma}\right)\exp\left(-\Phi_{ix}\exp\left(-\frac{u}{\nu} - \overline{\gamma}\right)\right)\mathrm{d}u = -\nu\Phi_{ix}\int_0^{\infty}\left(\ln y + \overline{\gamma}\right)\exp\left(-\Phi_{ix}y\right)\mathrm{d}y.$$

Proceeding with the change of variables $z = \Phi_{ix}y$ yields:

$$\begin{aligned}
-\nu\Phi_{ix}\int_0^{\infty}\left(\ln y + \overline{\gamma}\right)\exp\left(-\Phi_{ix}y\right)\mathrm{d}y &= -\nu\int_0^{\infty}\left(\ln z - \ln\Phi_{ix} + \overline{\gamma}\right)\exp\left(-z\right)\mathrm{d}z \\
&= -\nu\int_0^{\infty}\exp\left(-z\right)\ln z\,\mathrm{d}z + \nu\left(\ln\Phi_{ix} - \overline{\gamma}\right).
\end{aligned}$$

Since

$$\int_0^{\infty}\exp\left(-z\right)\ln z\,\mathrm{d}z = \frac{\partial}{\partial w}\int_0^{\infty} z^w\exp\left(-z\right)\mathrm{d}z\Big|_{w=0} = \Gamma'(1) = -\overline{\gamma},$$

we obtain:

$$\mathbb{E}\left[\max_{i'}\left\{\beta V_{i'xt+1} - \tau_{ii'} + \nu\varphi_{i'xt}\right\}\right] = \nu\ln\Phi_{ix} = \nu\ln\sum_{i'=0}^{N}\exp\left(\frac{\beta V_{i'xt+1} - \tau_{ii'}}{\nu}\right).$$

Plugging into V_{ixt} yields:

$$V_{ixt} = \ln w_{ixt} + \nu \ln \sum_{i'=0}^{N} \exp\left(\frac{\beta V_{i'xt+1} - \tau_{ii'}}{\nu}\right).$$

Given log utility, the expected value admits the following money-metric representation:

$$\widetilde{V}_{ixt} = \exp\left(V_{ixt}\right) = w_{ixt}\left[\sum_{i'=0}^{N} \exp\left(-\frac{\tau_{ii'}}{\nu}\right) \widetilde{V}_{i'xt+1}^{\frac{\beta}{\nu}}\right]^{\nu}.$$

In steady state, we drop the time subscript to obtain (5).

Finally, let $\mu_{ii'xt}$ denote the fraction of type-x analysts switching jobs from i to i' at the end of time t, by the law of large numbers, we have:

$$\begin{aligned}
\mu_{ii'xt} &= \Pr\left(\arg\max_{k}\left\{\beta V_{kxt+1} - \tau_{ik} + \nu\varphi_{kxt}\right\} = i'\right) \\
&= \int_{-\infty}^{\infty} \prod_{k\neq i'} \Pr\left(\beta V_{kxt+1} - \tau_{ik} + \nu\varphi_{kxt} \leq u\right) \mathrm{d}\Pr\left(\beta V_{i'xt+1} - \tau_{ii'} + \nu\varphi_{i'xt} \leq u\right) \\
&= \int_{-\infty}^{\infty} \prod_{k\neq i'} F\left(\frac{u - \beta V_{kxt+1} + \tau_{ik}}{\nu}\right) \mathrm{d}F\left(\frac{u - \beta V_{i'xt+1} + \tau_{ii'}}{\nu}\right) \\
&= \frac{1}{\nu}\exp\left(\frac{\beta V_{i'xt+1} - \tau_{ii'}}{\nu}\right) \int_{-\infty}^{\infty} \exp\left(-\Phi_{ix}\exp\left(-\frac{u}{\nu} - \overline{\gamma}\right)\right) \exp\left(-\frac{u}{\nu} - \overline{\gamma}\right) \mathrm{d}u \\
&= \frac{\exp\left(-\frac{\tau_{ii'}}{\nu}\right) \widetilde{V}_{i'xt+1}^{\frac{\beta}{\nu}}}{\sum_{k=0}^{N} \exp\left(-\frac{\tau_{ik}}{\nu}\right) \widetilde{V}_{kxt+1}^{\frac{\beta}{\nu}}},
\end{aligned} \tag{28}$$

The law of motion of type-x analysts is therefore given by:

$$a_{ixt+1} = \sum_{i'=0}^{N} \mu_{i'ixt} a_{i'xt}.$$

Substituting $\mu_{i'ixt}$ using equation (28) and dropping the time subscript yields (6), where $\mu_{i'ix}$ corresponds to (7). □

B.3 Proof of Proposition 3

In steady state, conditional on agreement over the next-period payoff w_{ix}, equation (5) implies that the surplus of an infinitesimal mass $\epsilon \to 0$ of type-x analysts is given by:

$$\epsilon w_{ix}^{*} \left[\sum_{i'=0}^{N} \exp\left(-\frac{\tau_{ii'}}{\nu}\right) \widetilde{V}_{i'x}\left(w_{i'x}; \boldsymbol{w}_{\neg i'x}^{*}\right)^{\frac{\beta}{\nu}} \right]^{\nu},$$

where w_{ix}^{*} is the current-period equilibrium payoff preset in the prior period. The composite term $\left[\sum_{i'=0}^{N} \exp\left(-\frac{\tau_{ii'}}{\nu}\right) \widetilde{V}_{i'x}\left(w_{i'x}; \boldsymbol{w}_{\neg i'x}^{*}\right)^{\frac{\beta}{\nu}}\right]^{\nu}$ captures the next-period option value. Bargaining in the current period affects only the option value through $\widetilde{V}_{ix}\left(w_{ix}; \boldsymbol{w}_{\neg ix}^{*}\right)$. Under disagreement, because the option value associated with broker i is no longer available in the next period, the analysts' surplus becomes:

$$\epsilon w_{ix}^{*} \left[\sum_{i' \neq i} \exp\left(-\frac{\tau_{ii'}}{\nu}\right) \widetilde{V}_{i'x}\left(w_{i'x}; \boldsymbol{w}_{\neg i'x}^{*}\right)^{\frac{\beta}{\nu}} \right]^{\nu}.$$

Therefore, the analysts' surplus from agreement relative to disagreement is given by:

$$\epsilon w_{ix}^{*} \left\{ \left[\sum_{i'=0}^{N} \exp\left(-\frac{\tau_{ii'}}{\nu}\right) \widetilde{V}_{i'x}\left(w_{i'x}; \boldsymbol{w}_{\neg i'x}^{*}\right)^{\frac{\beta}{\nu}} \right]^{\nu} - \left[\sum_{i' \neq i} \exp\left(-\frac{\tau_{ii'}}{\nu}\right) \widetilde{V}_{i'x}\left(w_{i'x}; \boldsymbol{w}_{\neg i'x}^{*}\right)^{\frac{\beta}{\nu}} \right]^{\nu} \right\}. \tag{29}$$

Since both ϵ and w_{ix}^{*} are constant, only the term in the curly brackets is relevant for Nash bargaining. In equilibrium, this term corresponds to $S_{ix}^{A}\left(\boldsymbol{w}_{x}^{*}\right)$ in equation (9).

For broker i, conditional on agreement over the next-period payoff w_{ix} to type-x analysts, its surplus attributable to an infinitesimal mass $\epsilon \to 0$ of type-x analysts is given by:

$$\epsilon \left[\frac{\partial S_{i}^{T}\left(\boldsymbol{q}\left(\boldsymbol{W}\right)\right)}{\partial a_{ix}\left(w_{ix}; \boldsymbol{w}_{\neg ix}^{*}\right)} - w_{ix} \right].$$

Under disagreement, the mass ϵ of type-x analysts moves from broker i to other options $i' \neq i$ in the next period. Using the job-switching fractions given by equation (7), the corresponding change in broker i's surplus is:

$$\epsilon \sum_{i' \neq i} \frac{\exp\left(-\frac{\tau_{ii'}}{\nu}\right) \widetilde{V}_{i'x}\left(w_{i'x}; \boldsymbol{w}_{\neg i'x}^{*}\right)^{\frac{\beta}{\nu}}}{\sum_{k \neq i} \exp\left(-\frac{\tau_{ik}}{\nu}\right) \widetilde{V}_{kx}\left(w_{kx}; \boldsymbol{w}_{\neg kx}^{*}\right)^{\frac{\beta}{\nu}}} \frac{\partial S_{i}^{T}\left(\boldsymbol{q}\left(\boldsymbol{W}\right)\right)}{\partial a_{i'x}\left(w_{i'x}; \boldsymbol{w}_{\neg i'x}^{*}\right)}.$$

Therefore, the broker's surplus from agreement relative to disagreement is given by:

$$\epsilon \left[\frac{\partial S_i^T \left(\boldsymbol{q} \left(\boldsymbol{W} \right) \right)}{\partial a_{ix} \left(w_{ix}; \boldsymbol{w}_{\neg ix}^* \right)} - \sum_{i' \neq i} \frac{\exp\left(-\frac{\tau_{ii'}}{\nu} \right) \widetilde{V}_{i'x} \left(w_{i'x}; \boldsymbol{w}_{\neg i'x}^* \right)^{\frac{\beta}{\nu}}}{\sum_{k \neq i} \exp\left(-\frac{\tau_{ik}}{\nu} \right) \widetilde{V}_{kx} \left(w_{kx}; \boldsymbol{w}_{\neg kx}^* \right)^{\frac{\beta}{\nu}}} \frac{\partial S_i^T \left(\boldsymbol{q} \left(\boldsymbol{W} \right) \right)}{\partial a_{i'x} \left(w_{i'x}; \boldsymbol{w}_{\neg i'x}^* \right)} - w_{ix} \right]. \quad (30)$$

Since ϵ is constant, only the term in the square brackets is relevant for Nash bargaining. In equilibrium, this term corresponds to $S_{ix}^B \left(\boldsymbol{W}^* \right)$ in equation (10).

The Nash bargaining solution is given by:

$$w_{ix}^* = \arg\max_{w_{ix}} \left\{ S_{ix}^A \left(\boldsymbol{w}_x \right)^{\alpha_x} S_{ix}^B \left(\boldsymbol{W} \right)^{1-\alpha_x} \right\}.$$

The FOC is:

$$\alpha_x \frac{\frac{\partial S_{ix}^A \left(\boldsymbol{w}_x^* \right)}{\partial w_{ix}^*}}{S_{ix}^A \left(\boldsymbol{w}_x^* \right)} + \left(1 - \alpha_x \right) \frac{\frac{\partial S_{ix}^B \left(\boldsymbol{W}^* \right)}{\partial w_{ix}^*}}{S_{ix}^B \left(\boldsymbol{W}^* \right)} = 0. \quad (31)$$

By equations (5) and (29), we have:

$$\frac{\partial S_{ix}^A \left(\boldsymbol{w}_x^* \right)}{\partial w_{ix}^*} = \beta \left(w_{ix}^* \right)^{\frac{1-\nu}{\nu}} \widetilde{V}_{ix} \left(\boldsymbol{w}_x^* \right)^{\frac{\beta-1}{\nu}} \frac{\partial \widetilde{V}_{ix} \left(\boldsymbol{w}_x^* \right)}{\partial w_{ix}^*}. \quad (32)$$

By equation (30), we obtain:

$$\frac{\partial S_{ix}^B \left(\boldsymbol{W}^* \right)}{\partial w_{ix}^*} = \frac{\partial}{\partial w_{ix}^*} \left(\frac{\partial S_i^T \left(\boldsymbol{q} \left(\boldsymbol{W}^* \right) \right)}{\partial a_{ix} \left(\boldsymbol{w}_x^* \right)} \right) - 1. \quad (33)$$

Finally, substituting (32) and (33) into (31) and rearranging yields equation (8).

C LLM Prompt for Extracting Broker Revenue Items

This appendix reports the ChatGPT 5.5 prompt used to extract broker revenue items from financial filings, including SEC Form 10-K filings, annual reports, and other comparable regulatory or company disclosures. To ensure data quality, all LLM-generated outputs are audited and corrected by human researchers against the original filings.

```
You are a financial data extraction assistant. Extract broker-revenue
   variables from the attached financial filing for 2019 only.
```

```
Use ONLY the filing.
Do NOT use outside knowledge.
Do NOT infer or estimate missing figures.
If a figure is not explicitly reported, return NA.
Be conservative and avoid double-counting.

STEP 1: Determine entity type and eligibility
Classify the filer into ONE category:
(a) Broker-dealer/securities firm
(b) Broker-dealer with asset/wealth management
(c) Bank/bank holding company with brokerage or capital-markets
   divisions
(d) Insurance/insurance-led financial group with incidental brokerage,
   wealth, or securities activities
(e) Asset/wealth manager/alternative manager
(f) Investment fund/pooled investment vehicle/umbrella fund
(g) Other/unclear mixed structure

Then classify:
entity_eligibility_for_broker_revenue = eligible/
   ineligible_fund_vehicle/ineligible_nonbroker_financial/unclear

If the filer is a fund vehicle or another clearly ineligible entity,
   return NA for narrow and broad broker revenue and explain briefly in
    the caveat.

STEP 2: Choose one reporting scope
If the filing contains multiple statement sets (Group, Bank, Parent,
   subsidiary, etc.), use only ONE scope.
Prefer the highest-level consolidated financial statements.
Do NOT mix scopes.

STEP 3: Extract only revenue-side items
Use this order:
```

```
1. consolidated statement of operations/income statement
2. revenue notes for the same reporting scope
3. segment/division note only if no relevant consolidated revenue item
   exists

Do NOT use:
- expense items
- business-description text as a source of numbers
- non-GAAP tables
- pro forma figures
- forecasts/guidance

Definitions:
- total_revenue = the main top-line revenue figure for 2019 in the
   chosen scope; if "Net revenues" is the top line, use that
- narrow_broker_revenue_item = the single cleanest revenue-side item
   tied to brokerage execution activity
- broad_broker_revenue_items_included = broader broker/capital-markets
   revenue items, but only if they are separately reported and non-
   overlapping

Preferred narrow item order:
1. Brokerage
2. Commissions from sales and trading
3. Commissions
4. Brokerage commissions
5. Transaction fees
6. Note-level brokerage breakdown from the same scope
7. Otherwise NA

Broad items may include, if separately reported and non-overlapping:
- Brokerage
- Commissions from sales and trading
- Commissions
```

```
- Underwriting
- Investment banking
- Investment banking - capital markets
- Investment banking - advisory
- Principal transactions, net
- Trading revenues

Exclude:
- net premiums/insurance premiums
- investment income of fund vehicles
- net gains/losses on investments of fund vehicles
- advisory or asset-management fees unless clearly brokerage-related
- bank deposit sweep income
- net interest income unless clearly part of broker revenue
- other/miscellaneous income
- expense-side commissions
- segment totals if consolidated line items exist

Broad revenue sum rule:
Sum broad items only if all are distinct, non-overlapping, from the
   same reporting scope, and not subtotals containing one another.
Otherwise return NA.

Convert all reported amounts into plain numbers with no commas and no
   unit label.
Example: 1033379000

Return ONLY:

firm_type:
entity_eligibility_for_broker_revenue:
reporting_scope_used:
reporting_currency:
```

```
total_revenue:
total_revenue_page:

narrow_broker_revenue_item:
narrow_broker_revenue_amount:
narrow_broker_revenue_page:

broad_broker_revenue_items_included:
broad_broker_revenue_amount_or_NA:
broad_broker_revenue_page:

short_note_on_overlap_or_caveat:
```

References

ABADIE, A., S. ATHEY, G. W. IMBENS, AND J. M. WOOLDRIDGE (2023): "When should you adjust standard errors for clustering?" Quarterly Journal of Economics, 138, 1–35.

ABOWD, J. M., F. KRAMARZ, AND D. N. MARGOLIS (1999): "High wage workers and high wage firms," Econometrica, 67, 251–333.

AMZALLAG, A., C. GUAGLIANO, AND V. LO PASSO (2021): "MiFID II research unbundling: Assessing the impact on SMEs," ESMA Working Paper.

ARTUÇ, E., S. CHAUDHURI, AND J. MCLAREN (2010): "Trade shocks and labor adjustment: A structural empirical approach," American Economic Review, 100, 1008–1045.

ASQUITH, P., M. B. MIKHAIL, AND A. S. AU (2005): "Information content of equity analyst reports," Journal of Financial Economics, 75, 245–282.

ATKESON, A. AND A. BURSTEIN (2008): "Pricing-to-market, trade costs, and international relative prices," American Economic Review, 98, 1998–2031.

BAGGER, J., F. FONTAINE, F. POSTEL-VINAY, AND J.-M. ROBIN (2014): "Tenure, experience, human capital, and wages: A tractable equilibrium search model of wage dynamics," American Economic Review, 104, 1551–1596.

BAKER, A. C., D. F. LARCKER, AND C. C. WANG (2022): "How much should we trust staggered difference-in-differences estimates?" Journal of Financial Economics, 144, 370–395.

BARBER, B., R. LEHAVY, M. MCNICHOLS, AND B. TRUEMAN (2001): "Can investors profit from the prophets? Security analyst recommendations and stock returns," Journal of Finance, 56, 531–563.

BONHOMME, S., T. LAMADON, AND E. MANRESA (2019): "A distributional framework for matched employer employee data," Econometrica, 87, 699–739.

BRADLEY, D., J. CLARKE, S. LEE, AND C. ORNTHANALAI (2014): "Are analysts' recommendations informative? Intraday evidence on the impact of time stamp delays," Journal of Finance, 69, 645–673.

BRADLEY, D., S. GOKKAYA, AND X. LIU (2017): "Before an analyst becomes an analyst: Does industry experience matter?" Journal of Finance, 72, 751–792.

BRADSHAW, M. T., Y. ERTIMUR, AND P. C. O'BRIEN (2017): "Financial analysts and their contribution to well-functioning capital markets," Foundations and Trends in Accounting, 11, 119–191.

BROWN, L. D., A. C. CALL, M. B. CLEMENT, AND N. Y. SHARP (2015): "Inside the 'black box' of sell-side financial analysts," Journal of Accounting Research, 53, 1–47.

CAHUC, P., F. POSTEL-VINAY, AND J.-M. ROBIN (2006): "Wage bargaining with on-the-job search: Theory and evidence," Econometrica, 74, 323–364.

CALDWELL, S., I. HAEGELE, AND J. HEINING (2026): "Bargaining and inequality in the labor market," Quarterly Journal of Economics, 141, 315–371.

CALIENDO, L., M. DVORKIN, AND F. PARRO (2019): "Trade and labor market dynamics: General equilibrium analysis of the China trade shock," Econometrica, 87, 741–835.

CHANG, B. AND H. HONG (2017): "Assignment of stock market coverage," NBER Working Paper, 23115.

——— (2019): "Selection versus talent effects on firm value," Journal of Financial Economics, 133, 751–763.

CLEMENT, M. B. (1999): "Analyst forecast accuracy: Do ability, resources, and portfolio complexity matter?" Journal of Accounting and Economics, 27, 285–303.

CLEMENT, M. B., J. HALES, AND Y. XUE (2011): "Understanding analysts' use of stock returns and other analysts' revisions when forecasting earnings," Journal of Accounting and Economics, 51, 279–299.

CLEMENT, M. B., L. KOONCE, AND T. J. LOPEZ (2007): "The roles of task-specific forecasting experience and innate ability in understanding analyst forecasting performance," Journal of Accounting and Economics, 44, 378–398.

COLLARD-WEXLER, A., G. GOWRISANKARAN, AND R. S. LEE (2019): "'Nash-in-Nash' bargaining: A microfoundation for applied work," Journal of Political Economy, 127, 163–195.

CORREIA, S., P. GUIMARÃES, AND T. ZYLKIN (2020): "Fast Poisson estimation with high-dimensional fixed effects," The Stata Journal, 20, 95–115.

CRAWFORD, G. S., R. S. LEE, M. D. WHINSTON, AND A. YURUKOGLU (2018): "The welfare effects of vertical integration in multichannel television markets," Econometrica, 86, 891–954.

CULLEN, Z. (2024): "Is pay transparency good?" Journal of Economic Perspectives, 38, 153–180.

DERRIEN, F. AND A. KECSKÉS (2013): "The real effects of financial shocks: Evidence from exogenous changes in analyst coverage," Journal of Finance, 68, 1407–1440.

DIETHER, K. B., C. J. MALLOY, AND A. SCHERBINA (2002): "Differences of opinion and the cross section of stock returns," Journal of Finance, 57, 2113–2141.

DRAKE, M. S., J. R. MOON JR, B. J. TWEDT, AND J. D. WARREN (2023): "Social media analysts and sell-side analyst research," Review of Accounting Studies, 28, 385–420.

EATON, J., D. JINKINS, J. R. TYBOUT, AND D. XU (2026): "Two-sided search in international markets," Journal of Political Economy, Forthcoming.

EATON, J. AND S. KORTUM (2002): "Technology, geography, and trade," Econometrica, 70, 1741–1779.

EECKHOUT, J. AND P. KIRCHER (2011): "Identifying sorting—in theory," Review of Economic Studies, 78, 872–906.

EECKHOUT, J., R. PINHEIRO, AND K. SCHMIDHEINY (2014): "Spatial sorting," Journal of Political Economy, 122, 554–620.

EU (2024): "Directive (EU) 2024/2811 of the European Parliament and of the Council of 23 October 2024 amending Directive 2014/65/EU," Official Journal of the European Union, OJ L, 2024/2811.

FANG, B. AND O.-K. HOPE (2021): "Analyst teams," Review of Accounting Studies, 26, 425–467.

FCA (2019): "Implementing MiFID II: Multi-firm review of research unbundling reforms," Financial Conduct Authority Multi-Firm Review.

——— (2024): "Payment optionality for investment research," Financial Conduct Authority Policy Statement PS24/9.

——— (2025): "Investment research payment optionality for fund managers," Financial Conduct Authority Policy Statement PS25/4.

FUJITA, S., G. MOSCARINI, AND F. POSTEL-VINAY (2024): "Measuring employer-to-employer reallocation," American Economic Journal: Macroeconomics, 16, 1–51.

GAO, X. AND J. R. RITTER (2010): "The marketing of seasoned equity offerings," Journal of Financial Economics, 97, 33–52.

GARICANO, L. (2000): "Hierarchies and the organization of knowledge in production," Journal of Political Economy, 108, 874–904.

GLEASON, C. A. AND C. M. LEE (2003): "Analyst forecast revisions and market price discovery," The Accounting Review, 78, 193–225.

GOKKAYA, S., X. LIU, V. K. POOL, F. XIE, AND J. ZHANG (2023): "Is there investment value in the soft-dollar arrangement? Evidence from mutual funds," Review of Financial Studies, 36, 3122–3162.

GOWRISANKARAN, G., A. NEVO, AND R. TOWN (2015): "Mergers when prices are negotiated: Evidence from the hospital industry," American Economic Review, 105, 172–203.

GREEN, T. C., R. JAME, S. MARKOV, AND M. SUBASI (2014): "Access to management and the informativeness of analyst research," Journal of Financial Economics, 114, 239–255.

GROYSBERG, B., P. M. HEALY, AND D. A. MABER (2011): "What drives sell-side analyst compensation at high-status investment banks?" Journal of Accounting Research, 49, 969–1000.

GUO, Y. AND L. MOTA (2021): "Should information be sold separately? Evidence from MiFID II," Journal of Financial Economics, 142, 97–126.

HARFORD, J., F. JIANG, R. WANG, AND F. XIE (2019): "Analyst career concerns, effort allocation, and firms' information environment," Review of Financial Studies, 32, 2179–2224.

HERKENHOFF, K., J. LISE, G. MENZIO, AND G. M. PHILLIPS (2024): "Production and learning in teams," Econometrica, 92, 467–504.

HILARY, G. AND C. HSU (2013): "Analyst forecast consistency," Journal of Finance, 68, 271–297.

HONG, H. AND M. KACPERCZYK (2010): "Competition and bias," Quarterly Journal of Economics, 125, 1683–1725.

HONG, H. AND J. D. KUBIK (2003): "Analyzing the analysts: Career concerns and biased earnings forecasts," Journal of Finance, 58, 313–351.

HONG, H., J. D. KUBIK, AND A. SOLOMON (2000a): "Security analysts' career concerns and herding of earnings forecasts," RAND Journal of Economics, 121–144.

HONG, H., T. LIM, AND J. C. STEIN (2000b): "Bad news travels slowly: Size, analyst coverage, and the profitability of momentum strategies," Journal of Finance, 55, 265–295.

HUANG, A. H., A.-P. LIN, AND A. Y. ZANG (2022): "Cross-industry information sharing among colleagues and analyst research," Journal of Accounting and Economics, 74, 101496.

HUANG, A. H., A. Y. ZANG, AND R. ZHENG (2014): "Evidence on the information content of text in analyst reports," The Accounting Review, 89, 2151–2180.

JACKSON, A. R. (2005): "Trade generation, reputation, and sell-side analysts," Journal of Finance, 60, 673–717.

JÄGER, S., C. ROTH, N. ROUSSILLE, AND B. SCHOEFER (2024): "Worker beliefs about outside options," Quarterly Journal of Economics, 139, 1505–1556.

JEFFERS, J. S. (2024): "The impact of restricting labor mobility on corporate investment and entrepreneurship," Review of Financial Studies, 37, 1–44.

KAMBOUROV, G. AND I. MANOVSKII (2009): "Occupational mobility and wage inequality," Review of Economic Studies, 76, 731–759.

KELLY, B. AND A. LJUNGQVIST (2012): "Testing asymmetric-information asset pricing models," Review of Financial Studies, 25, 1366–1413.

KENT, R. (2023): "UK investment research review," HM Treasury Independent Review.

KLEINMAN, B., E. LIU, AND S. J. REDDING (2023): "Dynamic spatial general equilibrium," Econometrica, 91, 385–424.

KOTHARI, S. P., E. SO, AND R. VERDI (2016): "Analysts' forecasts and asset pricing: A survey," Annual Review of Financial Economics, 8, 197–219.

KUEHN, L.-A., M. SIMUTIN, AND J. J. WANG (2017): "A labor capital asset pricing model," Journal of Finance, 72, 2131–2178.

KUMAR, A., V. RANTALA, AND R. XU (2022): "Social learning and analyst behavior," Journal of Financial Economics, 143, 434–461.

LAMADON, T., M. MOGSTAD, AND B. SETZLER (2022): "Imperfect competition, compensating differentials, and rent sharing in the US labor market," American Economic Review, 112, 169–212.

LIND, N. AND N. RAMONDO (2023): "Trade with correlation," American Economic Review, 113, 317–353.

LISE, J. AND J.-M. ROBIN (2017): "The macrodynamics of sorting between workers and firms," American Economic Review, 107, 1104–1135.

LOH, R. K. AND R. M. STULZ (2011): "When are analyst recommendation changes influential?" Review of Financial Studies, 24, 593–627.

——— (2018): "Is sell-side research more valuable in bad times?" Journal of Finance, 73, 959–1013.

LYS, T. AND S. SOHN (1990): "The association between revisions of financial analysts' earnings forecasts and security-price changes," Journal of Accounting and Economics, 13, 341–363.

MARX, M., D. STRUMSKY, AND L. FLEMING (2009): "Mobility, skills, and the Michigan non-compete experiment," Management Science, 55, 875–889.

MENZIO, G. AND S. SHI (2011): "Efficient search on the job and the business cycle," Journal of Political Economy, 119, 468–510.

MO, H. (2024): "Cross-industry technological information spillovers through analyst networks," SSRN Working Paper, 5011339.

MOLLOY, R., R. TREZZI, C. L. SMITH, AND A. WOZNIAK (2016): "Understanding declining fluidity in the US labor market," Brookings Papers on Economic Activity, 2016, 183–259.

PAYNE, J. L. AND W. B. THOMAS (2003): "The implications of using stock-split adjusted I/B/E/S data in empirical research," The Accounting Review, 78, 1049–1067.

PHUA, K., M. THAM, AND C. WEI (2023): "Peer effects in equity research," Journal of Financial and Quantitative Analysis, 58, 647–676.

POSTEL-VINAY, F. AND J.-M. ROBIN (2002): "Equilibrium wage dispersion with worker and employer heterogeneity," Econometrica, 70, 2295–2350.

ROBERT HALF (2025): "Employment trends spotlight: Financial services," Robert Half Research and Insights.

ROCHET, J.-C. AND J. TIROLE (2003): "Platform competition in two-sided markets," Journal of the European Economic Association, 1, 990–1029.

SEC (2017): "Securities Industry and Financial Markets Association, SEC staff no-action letter," Securities and Exchange Commission.

——— (2019): "Securities Industry and Financial Markets Association, SEC staff no-action letter (extension)," Securities and Exchange Commission.

SIFMA (2023): "SIFMA statement on SEC not extending MiFID II no-action relief," Securities Industry and Financial Markets Association Press Release.

SONNEY, F. (2009): "Financial analysts' performance: Sector versus country specialization," Review of Financial Studies, 22, 2087–2131.

STARR, E. P., J. J. PRESCOTT, AND N. D. BISHARA (2021): "Noncompete agreements in the US labor force," Journal of Law and Economics, 64, 53–84.

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

UYEDA, M. T. (2023): "Statement on the expiration of the SEC staff no-action letter re: MiFID II," SEC Statement.

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