

Swiss Finance Institute Research Paper Series N°20-84

Artificial Intelligence and High-Skilled Work:
Evidence from Analysts



Jillian Grennan

Duke University

Roni Michaely

University of Geneva and Swiss Finance Institute

Artificial Intelligence and High-Skilled Work: Evidence from Analysts*

Jillian Grennan

Roni Michaely

Duke University

University of Geneva and SFI

Fuqua School of Business

Geneva Finance Research Institute

August 26, 2020

Abstract

Policymakers fear artificial intelligence (AI) will disrupt labor markets, especially for high-skilled workers. We investigate this concern using novel, task-specific data for security analysts. Exploiting variation in AI's power across stocks, we show analysts with portfolios that are more exposed to AI are more likely to reallocate efforts to soft skills, shift coverage towards low AI stocks, and even leave the profession. Analyst departures disproportionately occur among highly accurate analysts, leaving for non-research jobs. Reallocating efforts toward tasks that rely on social skills improve consensus forecasts. However, increased exposure to AI reduces the novelty in analysts' research which reduces compensation.

JEL classification: G17, G24, J23, J24, J31, O33

Keywords: artificial intelligence, big data, technology, automation, sell-side analysts, job displacement, labor and finance, social skills, non-cognitive skills, tasks, skill premium, skill-biased technological change, compensation

* Authors: Grennan, Duke University (e-mail: jillian.grennan@duke.edu); Michaely, University of Geneva (e-mail: roni.michaely@unige.ch). We thank Jason Furman, William Kerr, Joseph Pacelli, and seminar participants at the NBER Economics of AI Conference and NBER Personnel Economics for helpful comments on an earlier draft of this study. We thank Randall Dalton and Lin Zhao for excellent research assistance. Some of the data used in this study comes from TipRanks, a firm in which Roni Michaely has a stock interest and serves on the Board of Directors.

Artificial intelligence (AI) is a powerful form of automation that programs machines to act more like humans and is considered the most important general-purpose technology of this era (Brynjolfsson and McAfee, 2014; Goldfarb et al., 2020). AI is expected to impact nearly all aspects of society. To name just a few, AI promises to personalize medicine, enhance security, improve transportation, and make education more effective (Kratsios and Parker, 2020; Castellanos, 2020). Consistent with the potential for disruptive growth, last year's private investment in AI exceeded \$70 billion, an incredible 48% average annual growth rate (Mishra, 2020). What makes AI unique from other technologies like software and industrial robots is that it is predicted to displace high-skill occupations, especially those involving high levels of education and accumulated experience (Frank et al., 2019; Webb, 2020). This contrasts with a rich literature on technological changes that documents how they increase the relative wages of educated workers by complementing their skills (Katz and Murphy, 1992; Berman et al., 1994; Autor et al., 2003).

To determine in detail the impact of AI on high-skilled workers, we study sell-side analysts. These high-skilled workers are an ideal candidate to study because we have very detailed data on what they do on the job including tasks that require hard and soft skills, and evaluations of their performance, compensation, and product quality. This allows us to explore hypotheses that no other studies have explored. Importantly, in this context, there is also a clear definition of what AI is being used for. Part of the challenge facing any research study trying to analyze AI has been that no single definition exists. For example, saying that AI involves programming computers to do things which if done by humans would require intelligence can be construed in many ways. Yet for analysts, we know that AI is being used to make investment recommendations and predict earnings. Further, the tension between analysts and AI is similar to those for suppliers of information in other contexts such as those that forecast supply and demand or predict creditworthiness. Finally, our paper is able to encompass not only firm-specific investment in AI but also investment that

is external such as that coming from entrepreneurial startups that compete with the firm and challenge high-skilled workers.

From a theoretical perspective, we rely on the canonical task-based framework to hypothesize about how AI affects analysts (Autor, Levy, and Murnane, 2003; Acemoglu and Autor, 2011; Brynjolfsson, Mitchell, and Rock, 2018). If analysts' only task is prediction, the framework suggests that AI directly substitutes for analysts resulting in job loss. If analysts' jobs consist of multiple tasks, then this framework suggests that removing the prediction task from analysts' jobs allows them to focus their attention on complementary tasks. In this sense, AI increases analysts' productivity, for example, by allowing them to cover more stocks or engage with more institutional investors (Acemoglu and Restrepo, 2018a; Agrawal, Gans, and Goldfarb, 2019).

We construct a dataset at the stock-analyst-quarter level to explore the implications of AI for analysts' jobs, including their employment status, stocks that they cover, time spent on tasks requiring different skills, the quality of their work output, and indirectly their compensation. We gather novel data and introduce new proxies for analysts' time spent on tasks requiring hard vs. soft skills.¹ The first proxy applies natural language processing (NLP) techniques to earnings call transcripts to quantify the complexity and content of analysts' questions. The second proxy stems from the collection of novel data on the number of meetings analysts have with management and institutional investors.

To measure AI intensity for a given stock, we quantify the amount of social media data that the AI algorithms process for a given stock. Our data is the same data used in a popular AI investment tool from TipRanks, which provides users with a "Smart Score," an AI generated rating that provides an indication of whether the stock will outperform based on eight factors extracted from unique datasets. Such a proxy for AI intensity matters to

¹Soft skills reflect high-cognitive skills such as coordination skills that complement knowledge-intensive, collaborative, networked activities like creative activities but also those that require a capacity to attribute mental states to others based on their behavior (Deming, 2017; Deming and Kahn, 2018).

analysts even if they do not pay attention to social media. A growing body of empirical evidence shows that there is “crowd wisdom” in social media data and that is incorporated into stock prices (Da et al., 2020; Da and Huang, 2020; Grennan and Michaely, 2020). Thus, even if analysts do not pay attention to social media or even if analysts’ firms do not use AI tools, they will feel competitive pressure from AI because it will be harder for them to bring value to their clients through the traditional means of selling research reports that contain novel trading angles.

Importantly, the AI intensity proxy from social media data has meaningful variation across stocks and sectors. To put the variation in perspective, Apple, Facebook, and Tesla get disproportionate social media attention: 100 times that of household names like Starbucks or Coca-Cola and 1000 times that of non-household names such as regional manufacturers or younger biotechs. Intuitively, our empirical analyses makes use of the variation that stems from some analysts having more stocks in their portfolio classified as high AI stocks, which are defined by being in the highest quartile of social media posts. In our sample, the average analyst covers 8.5 stocks and 2.0 stocks are classified as high AI intensity. A one standard deviation increase is associated with an analyst covering 4.4 stocks classified as high AI stocks, or roughly half of the stocks in their portfolio.

We begin our analysis by examining the potential for AI to displace analysts. We run ordinary least squares (OLS) regressions and find a positive correlation between the percent of stocks that analysts cover with high AI intensity and analysts’ decisions to leave the profession. Further, we see that analysts who remain employed as analysts are more likely to drop their coverage of stocks with higher AI intensity as well as initiate new coverage on stocks with lower AI intensity. These conditional correlations are consistent with AI substituting for high-skilled work, but they could simply be the result of endogeneity concerns.

To understand the full endogeneity concerns, it is worth taking seriously the set of assumptions that might make OLS estimates valid. For example, if analyst portfolios are

random with respect to AI, then the OLS would be informative. But for example, if AI intensity is associated with simple, uniform products with persistent demand, then, the less complex information environment may necessitate lower analyst coverage. In such a case, high AI intensity would be negatively correlated with analysts' departures and failures to initiate coverage, suggesting the OLS estimate understates the true effect. On the other hand, it is possible that AI intensity is associated with unexpected news content. If the unexpected occurs randomly, then OLS may still be informative. Whereas if the unexpected correlates with analysts' preferred firms to cover, the OLS estimates may be positively biased. Overall, the direction of the bias is ambiguous but more likely to be understated.

We introduce an instrumental variable (IV) to mitigate endogeneity concerns. The strategy exploits the variation in newspaper headline length generated by editorial space limits predetermined by the number of advertisements sold. Headlines must adhere to an assigned column width and editors are responsible for generating a title that fits within the space limits (Robinson, 2019). Our IV is an indicator for a stock having headline lengths below the median headline length in the USA Today in a given quarter. The IV has a strong first stage. The intuition for why the IV is relevant is that social media users often create content by mimicking what is popular and simply add their own personal commentary (Detweiler, 2019). Short titles are an example of text designed to entice readers and induce popularity (i.e., click-bait).² Yet consistent with our key identifying assumption that headline length is close to random, we find headline length is uncorrelated with firm characteristics and news content for millions of articles. For example, even for firms reporting unexpected earnings surprises, we see no statistically significant difference in headline length.

First, our IV estimates show that analysts are leaving the profession at significantly higher rates when the stocks they cover have higher AI intensity. In terms of economic magnitude,

²Recent research in psychology that uses randomized trials has shown that internet users are more likely to click on an article or a search engine result when the title is short. In fact, potential readers click on short titles over links with longer titles even when the information content is exactly the same (Konnikova, 2014).

the baseline rate for leaving the profession is 3.9% per analyst-quarter and our estimates suggest a 15% increase over the baseline. The IV estimates go in the same direction as the OLS estimates, but the magnitude is 1.9x higher, which seems plausible. The IV estimates for highly accurate analysts leaving the profession are even more pronounced. While the baseline rate for leaving the profession is only 0.7%, our estimates suggest a 60% increase over the baseline.

Second, we test whether analysts are departing at higher rates than in previous periods by gathering resume data from LinkedIn searches for departures before and after increased exposure to AI. We find that analysts are departing at higher rates during our sample than in the years prior to our sample. We also find that analysts are increasingly departing for non-research roles relative to the years prior to our sample. Consistent with this transition away from tasks being automated by AI, the analysts with greater exposure to high AI stocks prior to leaving are significantly more likely to go to a non-research job. For example, these analysts are more likely to take a job in investor relations or jobs requiring strong communication and social skills. Third, our IV estimates show that the analysts who continue to be analysts significantly shift their coverage toward low AI stocks. The relationship is symmetric. Analysts significantly decrease coverage of high AI stocks (10% more than baseline) and fail to initiate coverage on high AI stocks (7% less than baseline).

Taken together these results suggest AI and the future of high-skilled work is likely to involve some direct substitution via displacement. Moreover, these results hold across a variety of specifications including: (i) across employer and analyst's main sector of coverage, (ii) using alternative constructions of AI intensity in the analyst's portfolio, (iii) using different perturbations of the definition of accuracy and AI intensity, (iv) in placebo tests that use stocks unlikely to be impacted by the instrument, and (v) in placebo tests using randomly assigned AI intensity.

In the second part of the paper, we explore an interesting nuance in this shift in stock

coverage. While the shift in stock coverage could indicate AI is a direct substitute for analysts it could also represent a more nuanced form of complementarity. For example, complementarity would suggest that this shift in coverage occurs because the enhanced predictive capabilities brought about by AI allows analysts to focus their efforts on the stocks where the information environment is more complex. For such stocks, meetings with management, their suppliers, and hosting industry conferences are the best methods to gather soft information. Of course, participating in such events requires analysts to use their interpersonal skills, something that cannot be automated. Beyond soft information, analysts also have soft skills, such as devoting time to institutional clients to provide color for their recommendations.

To disentangle the direct substitution hypothesis from the more nuanced complementarity hypothesis, we evaluate analysts' time spent on such tasks. We find strong support for the complementary tasks hypothesis. We evaluate analysts' efforts on earnings calls and find that analysts ask significantly more questions and more complex questions when AI intensity is higher. In addition, the content of the analysts' questions change. Analysts significantly reduce their questions about easy-to-measure topics like sales and profits, but they increase their questions about harder-to-measure topics like brand and engagement. We also examine analysts' participation in meetings with management and institutional investors and find they significantly increase their attendance at such events when the stocks they cover have higher AI intensity (25% above baseline).

Next, we find that the quality of analysts' work improves significantly. Analysts' accuracy improves and their bias decreases. Consensus forecasts also improve. The evidence for improvement is consistent with the complementarity hypothesis. We also find evidence to suggest that analysts exert more effort. Specifically, we consider the boldness of analysts' forecasts, which requires more effort given that analysts are typically reluctant to deviate from the consensus (Diether et al., 2002; Michaely et al., 2018), and we find that they issue

bolder forecasts.

Finally, given that the market reaction to an analyst's recommendation is an indication of its value, we examine the market's response to changes in analysts' recommendations to indirectly infer how their compensation may be changing in response to AI. The value of such a recommendation directly affects their compensation, their ranking by among institutional investors, and how institutions value them. This, in turn, affects how much trading volume the investors direct toward their brokerage house, which also factors into their compensation. We see that the market is less responsive to analysts' recommendations when AI intensity is high. This is holds for both price movement and trading volume, which suggests analysts are being compensated less.

The key takeaways of our paper is that the relation between AI and the future of high-skilled work is likely to involve some direct substitution but also greater incentives for analysts to use their soft skills – ones that cannot be automated – to engage with management and clients and to develop more complex investment ideas. While, thus far, this has been associated with a higher quality work product and effort, one can imagine that this is only the beginning of more rapid period of change as the decrease in pay and talent push the industry into a new equilibrium.

Our study contributes to the literature on labor market responses to technological change. We provide detailed micro-level evidence of high-skilled workers job displacement, improved output, and a shift by the labor-force to using alternative skills that are harder to replicate by technology. This complements recent evidence from industrial robots (Graetz and Michaels, 2018; Acemoglu and Restrepo, 2020; Dauth et al., 2018; Humlum, 2020) supporting displacement for low-skill workers. While we cannot speak to aggregate employment changes from technological pass-through effects (Aghion et al., 2020) or sectoral shifts (Bena and Simintzi, 2018; Zator, 2020), we do document that analysts leave for non-research jobs. This is consistent with gains to high-skilled workers from technological change being concentrated

in super star firms that produce the technology (Autor et al., 2020; Humlum and Meyer, 2020; Babina et al., 2020). Overall, our findings support the notion that AI has very different implications (Acemoglu and Restrepo, 2018c,b, 2019) than previous waves of technological change, especially for wages (Katz and Murphy, 1992) and employment (Autor et al., 2003; Autor and Dorn, 2013; Autor et al., 2013; Michaels et al., 2014).

We also contribute to rich literature in personnel economics (Bartel et al., 2007; Lazear and Shaw, 2007). We find that analysts direct efforts to soft skills, consistent with research showing increased returns for non-cognitive skills (Heckman and Kautz, 2012; Castex and Dechter, 2014; Deming, 2017; Deming and Weidmann, 2020). Our finding that compensation decreases as reliance on social skills increases informs studies on wages and cognitive occupations over time (Kambourov and Manovskii, 2009; Beaudry et al., 2014; Altonji et al., 2016) and employers reliance on social screens (Altonji and Pierret, 2001; Hoffmann et al., 2018; Erel et al., 2018; Hacamo and Kleiner, 2020). Our findings are also consistent with studies showing that improvements in IT shifted high-skilled workers into flexible, team-based workplace settings (Bresnahan, 1999; Autor et al., 2002; Bresnahan et al., 2002) and the need for an adaptable culture to support those shifts (Graham et al., 2018, 2019; Grennan, 2019).

We contribute to a growing innovation literature on AI's economic consequences (Brynjolfsson, Mitchell, and Rock, 2018; Agrawal, Gans, and Goldfarb, 2018; Erel, Stern, Tan, and Weisbach, 2018; Li, Raymond, and Bergman, 2020) and societal impact (Frey and Osborne, 2017; Felten et al., 2018; Arntz et al., 2017; Felten et al., 2019). Unpacking the details of how AI works is important for debates surrounding AI's role in inequality (Aghion et al., 2019; Webb, 2020), firm dynamics (Alekseeva et al., 2019; Bessen et al., 2019; Babina and Howell, 2020), and regulation (Lee et al., 2019). Finally, we contribute to work studying the performance of man vs. machines in finance (D'Acunto et al., 2019; Abis, 2020; Coleman et al., 2020), analysts' careers (Hong and Kubik, 2003; Merkley et al., 2017; Kempf, 2020), and analysts' behaviors (Barber et al., 2001; Malmendier and Shanthikumar, 2014).

I. Data

Our goal is to study the relationship between AI that produces analyst-like forecasts and analysts' labor market outcomes over time. To achieve this goal, we combine analyst and stock data to produce a panel data set at the analyst-stock-quarter level. [Table I](#) provides descriptive statistics about the data analyzed in this study. We winsorize all variables at the 1st and 99th percentiles to minimize the influence of outliers. A brief description of each of the main components of the dataset are described below. Detailed formulas for the variables are included in [Appendix A](#). Due to the merging of various datasets, our main sample period runs from 2010Q1 to 2016Q4.

A. *Analysts*

The data set on analysts is sourced from IBES, which follows analysts across employers. Our main measures of interests are indicators for when analysts' leave the profession, their stocks covered, and reporting quality. IBES no longer provides the name of the analyst or their employer in the earnings forecast file but they do in their recommendation file (Merkley et al., 2020). We were able to identify the analysts using a file from 2015 with the historic cross links. We use these names to identify All-star analysts according to Institutional Investor magazine. We also use these names to search for a random sample of 25 analysts per year that leave the profession and classify their next job as a research job (i.e., buy-side or asset management), a non-research job (i.e., finance, consulting, strategy, investor relations, same employer, different job), or as a FinTech job, entrepreneur, or other.

We calculate three quality measures for analysts' reports: (i) accuracy, (ii) bias, and (iii) boldness. Intuitively, accuracy is the absolute forecast error and bias is the signed forecast error. Boldness reflects distance from the consensus forecast. Analysts tend to issue forecasts that are conservative relative to the consensus, and such forecasts require less effort.

Boldness measures the extra effort analysts are willing to put into their forecasts as well as the degree of independence of analysts (Clement and Tse, 2005). We calculate forecast errors for each analyst-stock forecast as the difference between the consensus earnings per share (EPS) forecast and actual EPS, scaled by either the absolute value of the consensus EPS forecast or the stock price at the end of the previous quarter. We use the consensus provided by Compustat due to data issues with IBES (Ljnuqvist et al., 2009). We follow the prior literature and exclude firms with absolute consensus bias of less than \$0.10 per share from our analysis to avoid issues with small numbers.

There is a rich literature showing that analysts' career and coverage decisions are associated with characteristics of the stock as well as characteristics idiosyncratic to the analyst. As such, our main control variables include those that are standard in the literature: analyst coverage, firm size, daily return volatility, mean monthly returns, market-to-book ratio, volatility of return on stock (ROE), profitability, and if the stock is a member of the S&P 500. All of these controls match those used by Hong and Kacperczyk (2010). The typical analyst controls include general experience, stocks covered, industries covered, employer, and reputation as proxied by being nominated as an All-star analyst by Institutional Investor magazine. These reflect insights from a broad literature showing the importance of analysts' career paths (Mikhail et al., 1999; Hong and Kubik, 2003; Groysberg et al., 2011), potential conflicts of interest (Michaely and Womack, 1999; Agrawal and Chen, 2008) and the mitigating role of institutions (Ljungqvist et al., 2007), herding (Clement and Tse, 2005), reputation (Fang and Yasuda, 2009), communication with insiders (Cohen et al., 2010), and industry expertise (Kadan et al., 2012).

B. Soft Skills

Two important soft skill advantages that analysts have is their communication skills and their creativity. Thus, meetings with management, company employees, suppliers and

customers provide an opportunity for analyst to use their competitive advantage over AI to uncover soft information. We proxy for these soft skills in two ways. First, we obtain data on analyst’s participation in meetings with management and institutional investors from Capital IQ’s public companies event data set. The database includes information on meeting type, attendees, date, and host. We then create simple counts of the numbers of meetings analyst hold in relation to a specific stock in a given quarter. Second, we use Capital IQ’s earnings transcript data to proxy for analysts’ ability to understand complex concepts. Specifically, we use NLP to go through earnings call transcripts and determine the number of questions analysts ask, the complexity of the question, and the topic of the question. We classify the question topics into easy-to-measure and harder-to-measure based on the occurrence of key words.

C. AI-based Stock Predictions

Hundreds of stock market intelligence FinTechs synthesize many data sources, including nontraditional ones, to inform investment decisions (Grennan and Michaely, 2020). TipRanks, a popular FinTech in this space, provides an investing tool that allows users to see a “Smart Score” which is an AI-generated indication of whether the stock will outperform or underperform based on eight factors extracted from their unique datasets. Appendix [Figure B1](#) shows an example of the AI-generated analysis for Facebook. We use TipRanks social media data, specifically, the number of social media posts that can be linked back to a particular stock as our proxy for the degree of AI-based competition or AI intensity for a given quarter. To understand the penetration of AI among all the stocks that the analyst covers, we define stocks covered with high AI intensity as the percent of stocks in the analyst portfolio that are in the top quartile of social media coverage in a given quarter.

II. Empirical Strategy

Our measure of AI intensity varies at the stock level. Thus, endogeneity concerns relate to reasons why analysts' portfolios of stocks may not be random with respect to AI. For example, if AI intensity is associated with less complex information environments, fewer analysts may cover those stocks. On the other hand, if AI intensity is associated with a particular sector or type of news content, then certain analysts may be disproportionately exposed to AI through those channels. Overall, the direction of the bias is ambiguous but more likely to be understated.

The main source of identification in this study is the occurrence of short headlines in the USA Today newspaper that attract the attention of those that post to social media. We focus on the USA Today, because it is a high readership newspaper that is more likely to be read by social media users than the financial press. Headline length varies meaningfully but the variation is quasi-random because it is due to editorial space constraints set by advertisement sales (Robinson, 2019). The intuition for why headline length is relevant is that bloggers and social media users often create content by mimicking what is popular and simply add their own personal commentary (Detweiler, 2019). Short titles are an example of text designed to entice readers and induce popularity.³ We also conduct tests to support the exclusion restriction but first we want to establish the intuition for the IV.

To illustrate our IV, an example of a short Apple headline in the USA Today is “Apple unveils iPad Mini, new Macs,” and a long headline is “Apple CEO Cook mum on new products, says ‘we have more game changers in us.’” The USA Today often also publishes general interest articles such as “What exactly is Apple Music anyway?” The mean headline length in the USA Today is 51 characters and the interquartile range is 15 characters. As

³Recent research in psychology that uses randomized trials has shown that internet users are more likely to click on an article or a search engine result when the title is short. In fact, potential readers click on short titles over links with longer titles even when the information content is exactly the same (Konnikova, 2014).

a second illustration, we display the headlines for the Wells Fargo's scandal and settlement with the Consumer Financial Protection Bureau from September 9, 2016. As [Table B1](#) shows the headline length varies from 27 to 111 characters. That is, holding the content of the news constant, we see a wide variation in headline lengths, suggesting editorial space constraints rather than the importance of the news dictate their length.

If the assumption that headline length is uncorrelated with other drivers of analysts' career decisions, stock coverage, and reporting quality is valid, conditional on controls, then the IV strategy generates a consistent estimate. While it is possible that headline length might be correlated with stock characteristics or a news event that may make an analyst pay attention to the stock, we test the idea that headline length is uncorrelated with news content or importance and find that they are either not systematically correlated, or the magnitude is so small to be consequential. For example, consistent with the identifying assumption, [Table B2](#) shows that the IV is uncorrelated with variation in stock characteristics such as momentum, market-to-book ratio, profitability, ROE, and firm size. We evaluate over 7 million headlines over all newspapers in our sample and find no variable associated with firm characteristics is statistically significant at the 95th percentile. Moreover, the R-squared is only 0.10%. We do find evidence that newspaper fixed effects are significant, suggesting that the editor plays an important role.

A potential limitation of our IV assumption is if the USA Today editor exhibits selection bias in that she picks a shorter headline in order shift readers' attention, including analysts, to more important firm news making them want to cover the stock. To explore this possibility, we use Capital IQ's Key Developments database, which provides summaries of material news and events that may affect the market value of securities. It monitors over 100 key development types including executive changes, M&A rumors, SEC inquiries, etc. Each key development item includes announcement date and type. We focus on value-relevant events and match the dates of the value relevant events to USA Today headlines about that firm

on that day. We then examine these 431,000 headlines to determine whether value-relevant news content is an important determinant of headline length.

Table B3 shows that there is no association between headline length and value-relevant events either in the cross-section or within-firm over time. Further, when we focus on earnings events as these likely represent the most important news for firms, and we add controls for positive and negative earnings surprises, we see no difference in headline length either in magnitude or statistical significance for such surprises. Finally, when we examine non-earnings key events such as payout announcements, targeting by activist investors, and other key non-earnings events (e.g., announcements of M&A deals), we again see little difference in headline length for these value-relevant events.

Finally, we employ the model selection technique of LASSO (Efron, Hastie, Johnstone, and Tibshirani, 2004) to determine whether particular headline words that predict headline length systematically convey something meaningful about the firm. **Table B4** shows the words selected by the variable selection model along with how much variation they explain. Inspecting the words reveals that they are not associated with content but with their own length. For example, the word “available” or “financial” are associated with longer headlines while the words “talk” and “china” are associated with shorter headlines.

We believe these tests provide further support to the assumption that the length of the headline is not related to news content or the complexity of the news environment, and thus, it plausibly satisfies the exclusion restriction assumption. We use the short headline IV in two stage least squares regressions that vary in their primary unit of observation: (i) analyst, (ii) analyst-stock, and (iii) stock level. In the first stage, we estimate the relationship between having below median headline length in a given quarter and AI intensity, which varies at the stock level. To aggregate from the stock level to the analyst-level, we calculate the percent of stocks an analyst covers that are in the highest quartile of social media posts and refer to this as stocks covered with high AI intensity. In this case, the IV is the percent of stocks

the analyst covers that have short headlines in a given quarter. Our IV approach uses the fitted values from the first stage to predict the outcome of interest. The second stage for the analyst level regressions are as follows:

$$Analyst_Leaves_{jt} = \alpha + \beta AI_Intensity_{jt} + \theta X_{jt} + \delta_t + \gamma_m + \rho_e + \epsilon_{jt} \quad (1)$$

where *Analyst_Leaves_{jt}* captures the analyst *j*'s decision to end her career as an analyst in quarter *t*. *AI_Intensity_{jt}* is our first-stage fitted value for stocks covered with high AI intensity. *X_{jt}* is a vector of analyst observables that includes stocks covered, industries covered, general experience, and reputation. δ_t is a quarter fixed effect, γ_m is a fixed effect for the analyst's main industry covered, ρ_e is a fixed effect for the analyst's employer, and ϵ_{jt} is the unobservable error component. Because an analyst can only leave at most once in her career, we do not include analyst fixed effects. When considering the hypothesis related to analysts' engaging in complementary tasks, we replace *Analyst_Leaves_{jt}* with *Analyst_Meetings_{jt}*. Given that the frequency of meetings varies over time, we also include an analyst fixed effect in those regressions.

The second stage for the analyst-stock level regressions are as follows:

$$Ends_Stock_Coverage_{ijt} = \alpha + \beta AI_Intensity_{ijt} + \theta X_{ijt} + \mu_j + \delta_t + \gamma_m + \rho_e + \epsilon_{ijt} \quad (2)$$

where *Ends_Stock_Coverage_{ijt}* captures the analyst *j*'s decision to stop covering stock *i* in quarter *t*. *AI_Intensity_{ijt}* is the fitted value for AI intensity, where it is defined as the quantity of social media posts associated with stock *i* in quarter *t*. *X_{ijt}* is a vector of analyst *j* observables as well as stock *i* observables including total analyst coverage, firm size, daily return volatility, mean monthly return, Amihud's illiquidity ratio, log market-to-book ratio, volatility of ROE, profitability, and an indicator for if the stock is a member of the S&P 500, total value relevant news events, and headline word controls for a given quarter *t*. μ_j

is an analyst fixed effect, δ_t is a year fixed effect, γ_m is a fixed effect for the analyst's main industry covered, ρ_e is a fixed effect for the analyst's employer, and ϵ_{ijt} is the unobservable error component. Because an analyst only ends coverage of a stock once in most, if not all, cases, we cannot include both analyst and stock fixed effects. We do, however, consider regressions with stock rather than analyst fixed effects. When we test the dependent variable for analyst initiating coverage we use lagged covariates to better reflect the typical timeline. When we evaluate hypotheses where the dependent variable relates to analyst's reporting quality (e.g., accuracy, bias, and boldness), we include analyst-by-stock fixed effects in our regressions.

Finally, the second stage for the stock level regressions are as follows:

$$Complexity_{it} = \alpha + \beta AI_Intensity_{it-1} + \theta X_{it-1} + \kappa_i + \delta_t + \omega_h + \epsilon_{it-1} \quad (3)$$

where $Complexity_{it}$ measures the overall complexity of analysts' questions about stock i in quarter t . $AI_Intensity_{it-1}$ is the fitted value AI intensity in the previous quarter. Matching our lag structure for initiating coverage, we use lagged explanatory and control variables for earnings call outcomes. This reflects the fact that earnings calls happen once per quarter, and thereby, any test for changes in behavior need to reflect such a timeline. X_{it-1} is a vector of stock i observables for quarter $t - 1$. δ_t is a year fixed effect, κ_i is a stock fixed effect, ω_h is an industry fixed effect for the stock, and ϵ_{it} is the unobservable error component. We also consider specifications without stock fixed effects to better understand cross-sectional variation.

III. Results

In this section, we explore whether analysts respond to the big data and AI-generated stock recommendations that are available for stocks. We start by examining the substitution hypothesis by considering analysts' decisions to switch careers and/or change stock coverage. Then, we evaluate the complementary tasks hypothesis by focusing on hard versus soft skills. And finally, we conclude by exploring implications for product quality and pay.

A. *Analysts Leaving the Profession*

The left panel of [Figure 1](#) displays the career decisions of analysts in relation to the portion of stocks that they cover with high AI intensity and the right panel the career decisions of the most accurate analysts, respectively. It illustrates that analysts are leaving the profession at a greater rate when more of the stocks they cover have high AI intensity. The line of best fit shows a significant positive linear relationship for all analysts and in particular, for the most accurate analysts. Each plot includes additional controls for analysts' workload, experience, reputation, and time.

In Panel A of [Table II](#) we report the associations between analysts quitting and leaving the profession and AI intensity based on the endogenous OLS regressions. In Columns 1 and 2 we report estimates that indicate that covering a greater portion of stocks with high AI intensity is positively and statistically significantly associated with analysts quitting and leaving the profession. This positive relationships are a conditional correlation after including controls for workload, reputation, experience, as well as fixed effects for year, main industry of stock coverage, and employer. The point estimate in Column 1 indicates that a one standard deviation increase in high AI intensity stocks in an analyst's portfolio is associated with a 0.015 standard deviation increase in that analyst quitting and leaving the profession.

Columns 3 and 4 of Panel A report results for the most accurate analysts and reveal

that the most accurate analysts are 0.024 standard deviations more likely to leave. Finally, Columns 5 and 6 report the results for Institutional Investor Magazine All-star analysts. In contrast to accurate analysts, the point estimate for Institutional Investor magazine All-star analysts leaving the profession is insignificant and close to 0. One potential explanation is that the All-star accolade more closely links with popularity than accuracy (Emery and Li, 2009). We consider this nuance in hard vs. soft skills more fully when we evaluate complementary tasks.

In Panel B of Table II, we report our results for analysts quitting and leaving the profession based on IV regressions. The point estimate in Column 1 shows a statistically significant 0.029 standard deviation increase in analysts quitting for a one standard deviation increase in AI intensity. To put the point estimate in perspective, the average analyst in our sample covers 8.5 stocks and 2.0 stocks have high AI intensity. A one standard deviation increase in AI intensity is equivalent to covering 2.4 more stocks with high AI intensity and is associated with 0.6% increase in the likelihood of quitting and leaving the profession. The baseline quit rate is 3.9% per analyst-quarter, suggesting a 15% increase over the baseline. The IV estimate is 1.9x that of the OLS estimate, which seems plausible even though the direction of the bias is ambiguous. Further, the results hold when we include employer fixed effects suggesting this is not the result of recent consolidation in the banking industry.

Consistent with the figures and OLS estimates, the IV estimates for accurate analysts' quitting are even more pronounced than for analysts' quitting. A one standard deviation increase in the AI intensity in an accurate analyst's portfolio is associated with a 0.054 standard deviations increase in quitting and leaving the profession. To put this number in perspective, the baseline quit rate for the most accurate analysts is 0.073% per analyst-quarter and this suggests an increase of 0.045%, or about a 60% increase. In contrast to accurate analysts, we cannot reject the null hypothesis that AI intensity is not associated with All-star analysts quitting and leaving the profession. The point estimates for All-star

analysts are small and close to 0. Finally, weak instrument tests suggest that our instrument is relevant, and thus, producing a consistent estimate.

Our finding of analysts, especially highly accurate analysts, quitting and leaving the profession is quite robust to decisions regarding sampling and regression specification. [Table III](#) consistently documents positive and significant quit rates for analysts for different cuts of the sample (Panel A and B) and definitions of AI intensity (Panel C) or accuracy thresholds (Panel D). For example, we have focused on the full sample of analyst quits so far, but analysts who only cover one or two stocks may already be preparing to quit so it may be possible that the two stocks they cover are both high AI stocks. In this case, the preparing to quit analyst may appear as an outlier with 100% AI intensity and distort the results. While we control for the total number of stocks covered in all of our regressions, we want to ensure the relationship we capture is not being driven by such outliers, so we exclude the analyst-quarter observations where the number of stocks are in either the top or bottom 10 percent of the distribution. As Panel A reports, the results show that the point estimates and statistical significance increase for all but All-star analysts, who remain close to 0. In unreported analysis, we also re-run our regressions using a randomly assigned AI intensity from a distribution with the same mean and standard deviation as the true variable. The IV fails weak instrument tests and the OLS results are insignificant, suggesting our findings are not driven by some unobserved factor.

Next, we explore how similar the analyst quit rates are to the years before AI-generated stock recommendations became popular. The top panel of [Figure 2](#) shows the binned scatterplot plot of quit rates for analysts and highly accurate analysts over time after controlling for workload, experience, reputation, and employer fixed effects. These plots illustrate a positive and significant increase in quit rates over time. Further, they reveal that the increase in departure rates coincides with the time period when meaningful technical advances occurred to enable AI applications such as stock picking. In addition, Appendix [Figure B2](#) presents

the scatter plot of the raw quarterly quit rate for analysts who leave the profession over time. The peak quit rate occurred at the height of the great financial crisis, but interestingly, the rates in more recent years are approaching those of the financial crisis. Further, the line of best fit shows a positive trend in quit rates over time even in the raw data.

To understand if the analysts are leaving for the same jobs that they always left for, we gather resume data from LinkedIn for a random sample of approximately 25 analysts per year. As reported in Appendix [Table B4](#), the three most popular next jobs for our sample of analysts are jobs in asset management (19.6%), other finance roles such with a corporation or in private stock (16.1%), and as buy-side analysts (15.8%). Less popular job moves include into investor relations, with the same employer but in a different role, entrepreneurship, consulting, corporate strategy, and to a FinTech. We categorize the job destinations into research jobs (i.e., asset management and buy-side research), non-research jobs, and FinTech/entrepreneurship jobs.

The bottom panel of [Figure 2](#) shows the binned scatterplot plot of research versus non-research jobs for analysts over time after controlling for workload, experience, and reputation. These plots suggests a meaningful change in terms of where analysts go. We see a significant increase in analysts going to non-research jobs and a significant decrease in analysts going to research jobs. Consistent with the aggregate figures, Appendix [Figure B4](#) shows plots for specific job titles. For example, we see more analysts going into investor relations and fewer analysts going into asset management.

In [Table IV](#), we examine if the job destinations vary between low-AI and high-AI quits. Column 1 shows that high AI-quits are significantly less likely to go to research jobs and relatedly, significantly more likely to go to non-research jobs. We observe no statistical difference in their propensity to go to FinTechs. The results are statistically significant at the 95th percentile and the instrument continues to pass weak instrument tests even for the small sample. This result is consistent with the competitive pressure from AI making

it harder for analysts to bring value to their clients through better research, and thereby, incentivizing these analysts to pursue careers that does not compete with AI.

In summary, we find meaningful evidence that AI is substituting for analysts. The analysts are quitting and leaving the profession at higher rates once the stocks that they cover become high AI intensity stocks. Moreover, the analysts are quitting their jobs more often than identical analysts who do not cover as many high AI intensity stocks. This phenomena is even more acute for the highly-skilled analysts. The analysts' quit rates are higher than in previous periods. The analysts are also more likely to move to non-research jobs than in previous period, especially those that quit who were highly exposed to the strength with which AI is processing data and automating parts of their job.

B. The Stocks Covered by Analysts

Next, we examine the hypothesis that AI is a substitute for high-skilled work for the analysts who continue to work as analysts by investigating if these analysts change the stocks that they make forecasts about. As reported in [Table V](#), we find evidence of a substitution effect even among the analysts who stay. Columns 1 and 2 show that when AI intensity increases for a stock, the analyst is less likely to initiate coverage of that stock and Columns 3 and 4 show that analysts are more likely to end coverage of the stock when AI intensity increases. Panel A reports the endogenous OLS estimates and Panel B reports the IV estimates. We apply a comprehensive fixed effect structure in both panels, including fixed effects for analyst, the analyst's main industry of coverage, and year. We also include a variety of controls for analyst and stock characteristics. The point estimates exhibit qualitative similarity, but for discussion purposes, we focus on the IV results in Panel B. Even at the more disaggregated analyst-quarter level, the instrument continues to pass weak instrument tests; for example, in the first stage estimation, the IV is significant at the 99th percentile and the F -statistic associated with the first stage exceeds necessarily critical

values for potential bias.

Column 1 of Panel B shows that a one standard deviation increase in AI intensity is associated with a 0.043 standard deviation decrease in initiating coverage. This finding is significant at the 99th percentile. To put this estimate in context, an analyst initiates coverage on a new stock in approximately one out of every four quarters. Given the skewed nature of social media data, a standard deviation increase would put a stock in the upper quartile of social media coverage. Thus, when a stock moves into the upper quartile of social media coverage, an analyst decreases his likelihood of initiating coverage on that stock by 1.9% in a given quarter, or about 7% below the baseline rate. In Column 2, stock fixed effects rather than analyst fixed are included. In this case, the point estimate is about 50% larger.

Column 3 of Panel B shows that a one standard deviation increase in AI intensity is associated with a 0.018 standard deviation increase in ending coverage of that stock. The unconditional average of ending stock coverage is 4.8% in a given quarter and this one standard deviation increase is equivalent to a 0.5% increase in the likelihood of ending coverage, or about 10% above the baseline rate. Moving from analyst to stock fixed effects as is reported in Column 4 suggests AI intensity is associated with a 0.090 standard deviation increase in ending coverage. Given that again we see larger economic magnitudes with the alternative fixed effects, we take the estimates that control for some time-invariant unobserved factor about the analyst to be our most conservative estimates. We note that we cannot include both stock and analyst fixed effects as there is only one because the dependent variable is equal to one at most one time in the life of analyst-stock pairing.

While a direct test of the exclusion restriction is not possible, we follow the prior empirical literature and implement placebo tests where, for reasons unrelated to our identification strategy, certain subsamples may or may not receive treatment and are therefore immune to the IV (Bound and Jaeger, 2000; Altonji et al., 2005; Angrist et al., 2010). Specifically,

we document that for a subset of microcap stocks where we expect short headlines to have no first stage effect on AI intensity, the short headline instrument also has no reduced form effect on initiating or ending stock coverage. By showing for this subsample that the first stage relationship (effect of the IV on treatment) and the reduced form relationship (direct effect of the IV on the dependent variable) no longer exist, this is reassuring evidence that the relationship estimated in the full sample is not driven by unobservables. These results are reported in Panel A of [Table VI](#).

Finally, we also explore the consequences of relaxing the exclusion restriction assumption that we rely upon for identification. Specifically, we estimate a bound for this potential bias by using recent advances in the econometrics literature for producing plausibly exogenous estimates (Conley et al., 2012; Imbens and Rubin, 2015; Kippersluis and Rietveld, 2018). Panel B of [Table VI](#) summarizes the plausibly exogenous IV estimates for initiating coverage in Columns 1 and 2 and for ending coverage in Columns 3 and 4. While the IV estimates are significant at the 99th percentile, the plausibly exogenous IV estimates are also significant, albeit at the 95th or 90th percentile. Notably, the plausibly exogenous IV estimates with the analyst fixed effects are more significant. Thus, these tests help to establish that we have robust evidence in support of analysts changing the stocks they cover in response to AI.

C. Complementary Tasks

The complementary tasks supposition argues that AI adoption will enable analysts to focus their efforts on their competitive advantage relative to AI. For example, analysts could focus on factors that determine stock performance for which less information is readily available or they could focus more on their soft skills such as conducting meetings with management to gather new insights or in marketing their research to institutional investors. Similarly, analysts may choose to focus their efforts on gathering information on stocks that is orthogonal to the insights AI produces. For example, AI may be good at analyzing images

of parking lots to improve the forecast accuracy for sales. On the other hand, AI may not be good at measuring qualitative concepts like customer service, brand influence, and the prospects of risky investments in new technologies.

Table VII explores the complementary tasks hypothesis in the context of earnings calls. Using our IV estimation strategy, Panel A shows a positive relationship between AI intensity and analysts efforts on earnings calls. The observed relationships are similar when we compare the within-stock estimates (Columns 1 and 3) and the cross-sectional estimates (Columns 2 and 4). Specifically, we find that analysts ask more questions and the complexity of their questions increases when AI intensity is higher. Panel B examines the content of the analysts' questions. Consistent with the complementary tasks hypothesis, we see analysts reduce their questions about easy-to-measure topics like sales and profits, but they increase their questions about harder-to-measure topics like brand and engagement. Overall, the evidence from earnings calls suggests that the ability of AI to streamline and automate predictions for easier-to-measure aspects of firm performance is enabling analysts to reallocate their time and efforts toward the exploration of more complex topics and toward gaining deeper insights for their investment clients.

Table VIII reports the results for a second setting in which analysts' soft skills are conjectured to complement AI. Specifically, we examine analyst's participation in meetings with management and institutional investors. The point estimate in Column 1 shows a statistically significant 0.17 standard deviation increase in analysts' meetings with management or investors for a one standard deviation increase in AI intensity. This represents an economically meaningful increase and suggests that analysts participate in 2.7 additional meetings with management or investors in a given quarter when a stock moves from low to high AI intensity. Our sample statistics show that, on average, about 85% of these meetings are with management. Column 2 reports the results when we include employer fixed effects, which is important since some banks may be more likely to host corporate events. Yet even with

employer fixed effects, we still see a 0.14 standard deviation increase. Next, we show that there are statistically significant increases in meetings regardless of it they are with institutional investors (Columns 3 and 4) or management (Columns 5 and 6). Importantly, these results without analyst fixed effects allow us to understand how a combination of new and remaining talent approach their tasks. In Appendix [Table B6](#), we report the results with analyst fixed effects which suggest that analysts focus relatively more on marketing their research to investors.

Overall, the tests from earnings calls and meetings support a task-based theory of how technology changes work. The tests suggest that AI serves as a replacement for easier prediction problems, but this then frees up time for employees to work on other tasks such as such as gathering soft information or selling research. This suggests at least in the context of analysts there is some complementarity from AI.

D. Product Quality and Effort

An interesting question that our data allows us to answer is how analysts' product quality relates to AI intensity. This question is important for two reasons. First, knowing the answer helps us to understand how analysts' efforts respond. Second, knowing the answer helps to provide a potential solution to the concern that delegating high-stakes decisions such as investment recommendations to a statistical model alone is problematic (Cowgill et al., 2020; Li et al., 2020). As more resources are delegated to AI systems, a potentially important safeguard is ensuring that any human predictions made alongside the AI recommendations remain of high quality. But given that AI represents a form of competition, analysts could strategically change the quality of their recommendations in the presence of AI (Lamont, 2002; Bond et al., 2012).

[Table IX](#) reports how the quality of the earnings forecasts relate to AI intensity. Using variation within analyst-stock pairs and our IV strategy, we find statistically significant

evidence to suggest that when AI intensity increases, analysts' accuracy increases and their bias decreases. This is a very important result because it suggests analysts are continuing to exert productive effort in the presence of AI. This estimate comes from a very tight regression specification that includes employer, main industry of coverage, and year fixed effects. The regression also includes controls for analyst and stock level controls, which help to account for other dynamics that could lead to a change in analyst reporting quality. The finding that analyst reporting quality improves is robust to our definition of reporting quality. As shown in Columns 3 and 4, when accuracy and bias are normalized by stock price as opposed to consensus estimate, the results are qualitatively similar and marginally statistically significant.

Further, the economic magnitude of the increase in reporting quality associated with this within analyst-stock pair specification is meaningful. The details of the point estimates, included in Appendix [Table B7](#), show the importance and convey additional useful information. First, AI intensity is not changing any traditional relationships observed in the data. For example, increased firm size and greater average monthly stock returns are both still associated with lower bias Hong and Kacperczyk (2010). Second, the results help put AI intensity in context by showing where it ranks relative to other controls. Our point estimates suggest that AI intensity has an economically meaningful effect on reporting quality, ranking near the middle of the point estimates reported. The point estimate for AI intensity is smaller than that of firm size and profitability but larger than that of monthly returns or analyst experience.

The final column of [Table IX](#) helps inform if analysts are exerting more effort. It presents evidence that suggests analysts are generating bolder forecasts. While bolder forecasts might generate more visibility and attention, they also require more effort, as analysts are typically reluctant to deviate from the consensus as the consequences of being wrong are high (Clement and Tse, 2005). That we find statistical significant evidence in support of greater effort is

important. For example, if reporting quality changes because of the talent pool for analysts, this would be difficult to reverse. In contrast, if reporting quality changes from low effort, contractual mechanisms could more easily be introduced. What is important is that we find no need for additional contractual mechanisms. In fact, it appears AI is enabling analysts to take advantage of their competitive advantage and engage in complementary tasks that enable them to ultimately produce bolder, and more accurate forecasts.

Next, we extend the product quality regressions to explore additional heterogeneity in the data. For example, Appendix [Table B8](#) shows the results from regressions with analyst-by-year fixed effects rather than following analyst-stock pairings over time. In contrast to the previous results, when we allow for heterogeneity in the stocks analysts cover, we find reporting quality declines in relation to AI intensity. This suggests that aggregate changes in product quality may not be the same as the within-stock estimates given that we know analysts are shifting their coverage away from high AI stocks.

To further explore this heterogeneity, we run quantile regressions for five quantiles of analyst accuracy (0.15, 0.25, 0.5, 0.75 and 0.85) and report those results in Appendix [Table B9](#). We observe a strictly monotonic relationship across the quantiles. Those in the lowest quantiles of accuracy produce less accurate research when AI intensity is high but this is not the case for those in the highest quantiles of accuracy. This again is an important result because it suggests that aggregate changes in product quality may not be the same as the within-stock estimates given that we know that the more accurate analysts are the ones leaving the profession.

As a final test of this heterogeneity, we return to our main specification with the analyst-by-stock fixed effects, but we introduce an indicator for analysts who are highly accurate historically and interact this term with AI intensity. As [Table B10](#) shows, we find a positive and significant interaction term, suggesting that the increase in accuracy is being driven by the highly accurate analysts. We also consider subsamples tests and find the same pattern.

Given that the improvements in reporting quality appear to be driven by the highly accurate analysts but the highly accurate analysts are also the ones leaving the profession at a higher rate, our last set of results evaluates overall accuracy and bias at the stock level. Thus, this allows for general equilibrium effects such as analysts leaving the profession or changing their stock coverage over time. As reported in [Table X](#), reporting quality is improving in relation to AI intensity, although the aggregate relationship is limited in part by these changes in the talent. We see that accuracy increases and bias decreases for analysts' forecasts both on average and at the median. Importantly, in Panel C, we evaluate consensus forecast error, and find that analysts' consensus forecast improves.

E. Implications for Analysts' Compensation

While there is no publicly available data set on analysts' compensation, we know their compensation is linked to abnormal trading volume generated by their research and to the overall value of their research (Brown et al., 2015). Thus, we examine the market's response to changes in analysts' recommendations to infer how their compensation may be changing in response to AI. We analyze excess returns and excess volume associated with recommendation changes. Specifically, we examine recommendation changes that involve an analyst upgrading a stock to a buy or a strong buy or downgrading a stock to a sell or a strong sell. For each recommendation, we estimate the cumulative abnormal returns (CARs) from the announcement. We use daily data to estimate the parameters of a Carhart four factor model in which the four factors are (1) the market return, which is the CRSP value-weighted index; (2) SMB (Small Minus Big), which is a mimicking portfolio to capture risk related to size; (3) HML (High Minus Low), which is a mimicking portfolio to capture risk associated with book-to-market ratio characteristics; and (4) UMD (Up Minus Down), which is a mimicking portfolio designed to address risk associated with prior returns by subtracting a portfolio of low prior return firms from a portfolio of high prior return firms. The event period is days 0

to +1. To align the signs correctly for downgrades, we multiply the CARs by -1. Abnormal volume is defined in a similar manner but by using the log transformed volume relative to a market model. Downgrades are not multiplied by -1 for volume.

Our specification uses the OLS estimates of the CARs along with our measure for AI intensity:

$$CAR_{ijt} = \alpha_{ijt} + \beta AI_Intensity_{ijt} + a_j + \delta_t + e_{ijt} \quad (4)$$

where β is the coefficient of interest, representing the market's relative change in response to analysts' recommendations as a function of $AI_Intensity_{ijt}$. To allow for analyst-specific effects, we include a_j . We also include δ_t , which captures time fixed effects.

Table XI presents the results of these tests of market responsiveness to analysts' recommendation revisions. The evidence suggests that the market is less responsive to analyst recommendations when AI intensity is high. The point estimates reported in Columns (1) and (2) of Panel A suggest a decrease in excess returns of around 24 to 27 basis point associated with a one standard deviation increase in AI intensity. These estimates are significant at the 99th percentile. The inclusion of analyst and time fixed effects ensures that these results are robust to factors affecting analyst recommendations such as general experience and all-star status as well as trends over time. Consistent with the results for excess returns, Columns (3) and (4) of Panel A show statistically significant decreases in excess trading volume when AI intensity of an stock is high. Both results indicate that when AI-generated stock picks are available, analysts' ability to find a novel and unique research angle is more difficult. In conclusion, we find indirect evidence to suggest that analysts' compensation is decreasing.

IV. Conclusion

AI is having a transformative effect on all types of industries and its applications are often performing some of the exact same tasks that highly-skilled workers do with greater accuracy. Thus, some commentators have argued that as more and more AI applications are deployed, labor markets will dramatically change. In this paper, we analyze the impact of AI in the context of security analysts.

Using a novel data and an IV approach, we find evidence that AI serves as direct substitute for analysts' work but also as a complement. As evidence of substitution, we find analysts leave the profession, especially highly accurate ones, while those who remain shift their coverage toward low-AI stocks. Analysts are departing at higher rates than in the past and they are increasingly leaving for non-research roles (e.g., to a job in investor relations).

Analysts' access to management gives them a soft information advantage over AI, and they increasingly focus their time on such meetings. Further, we find analysts spend more time on marketing (or on conveying less tangible information) with institutional investors. In addition to spending time on these tasks that require superior social skills, we also find support for analysts allocating their efforts toward more complex tasks. Specifically, we find evidence that analysts' questions on earnings call change: they ask more complex questions and shift their questions toward harder-to-measure rather than easier-to-measure concepts.

Finally, we find increased exposure to AI is associated with improved product quality: more accurate, less biased forecasts and improved consensus forecasts. In addition, we find evidence consistent with analysts exerting greater effort, perhaps from reallocating efforts toward complementary tasks that rely on social skills. Nevertheless, the novelty in analysts' research is lower, which we see in the market's reaction to their research, and hence, this provides evidence that their compensation is lower.

In conclusion, our findings are broadly consistent with the notion that a key consequence

of AI will be to shift high-skilled workers toward tasks that rely on social skills and complex thought. While those skills are certainly essential, they are often hard to improve in the short-term. This suggests that an important factor in how much AI will disrupt high-skilled work, and ultimately improve labor force productivity, is the extent to which high-skilled workers can be encouraged to invest in enhancing social skills.

REFERENCES

- Abis, S., 2020. Man vs. machine: Quantitative and discretionary equity management. working paper .
- Acemoglu, D., Autor, D., 2011. Handbook of Labor Economics, North-Holland, chap. Skills, Tasks, and Technologies: Implications for Employment and Earnings.
- Acemoglu, D., Restrepo, P., 2018a. Low-skill and high-skill automation. *Journal of Human Capital* 12, 204–232.
- Acemoglu, D., Restrepo, P., 2018b. Modeling automation. *AEA Papers and Proceedings* 108, 48–53.
- Acemoglu, D., Restrepo, P., 2018c. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review* 108, 1488–1542.
- Acemoglu, D., Restrepo, P., 2019. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives* 33, 3–30.
- Acemoglu, D., Restrepo, P., 2020. Robots and jobs: Evidence from US labor markets. *Journal of Political Economy* 128, 2188–2244.
- Aghion, P., Antonin, C., Bunel, S., Jaravel, X., 2020. What are the labor and product market effects of automation? New evidence from france. Tech. rep.
- Aghion, P., Jones, B., Jones, C., 2019. *The Economics of Artificial Intelligence*, University of Chicago Press, chap. Artificial Intelligence and Economic Growth.
- Agrawal, A., Chen, M., 2008. Do analyst conflicts matter? Evidence from stock recommendations. *Journal of Law and Economics* 51, 503–537.
- Agrawal, A., Gans, J. S., Goldfarb, A., 2018. Human judgment and ai pricing. *AEA Papers and Proceedings* 108, 58–63.

- Agrawal, A., Gans, J. S., Goldfarb, A., 2019. Artificial intelligence: The ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives* 33, 31–50.
- Alekseeva, L., Azar, J., Gine, M., Samila, S., Taska, B., 2019. The demand for ai skills in the labor market. working paper .
- Altonji, J. G., Elder, T. E., Taber, C. R., 2005. An evaluation of instrumental variable strategies for estimating the effects of catholic schooling. *Journal of Human Resources* 40, 791–821.
- Altonji, J. G., Kahn, L. B., Speer, J. D., 2016. Cashier or consultant? Entry labor market conditions, field of study, and career success. *Journal of Labor Economics* 34, S361–S401.
- Altonji, J. G., Pierret, C. R., 2001. Employer Learning and Statistical Discrimination. *Quarterly Journal of Economics* 116, 313–350.
- Angrist, J., Lavy, V., Schlosser, A., 2010. Multiple experiments for the causal link between the quantity and quality of children. *Journal of Labor Economics* 28, 773–824.
- Arntz, M., Gregory, T., Zierahn, U., 2017. Revisiting the risk of automation. *Economics Letters* 159, 157–160.
- Autor, D., Dorn, D., Katz, L. F., Patterson, C., Van Reenen, J., 2020. The Fall of the Labor Share and the Rise of Superstar Firms. *Quarterly Journal of Economics* 135, 645–709.
- Autor, D. H., Dorn, D., 2013. The growth of low-skill service jobs and the polarization of the u.s. labor market. *American Economic Review* 103, 1553–97.
- Autor, D. H., Dorn, D., Hanson, G. H., 2013. The china syndrome: Local labor market effects of import competition in the united states. *American Economic Review* 103, 2121–2168.
- Autor, D. H., Levy, F., Murnane, R. J., 2002. Upstairs, downstairs: Computers and skills on two floors of a large bank. *Industrial and Labor Relations Review* 55, 432–447.
- Autor, D. H., Levy, F., Murnane, R. J., 2003. The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics* 118, 1279–1333.
- Babina, T., Fedyk, A., He, A., Hodson, J., 2020. Artificial intelligence, firm growth, and industry concentration. working paper .
- Babina, T., Howell, S., 2020. Innovation investment and labor mobility: Employee entrepreneurship from corporate R&D. working paper .
- Barber, B., Lehavy, R., McNichols, M., Trueman, B., 2001. Can investors profit from the prophets? security analyst recommendations and stock returns. *Journal of Finance* 56, 531–563.

- Bartel, A., Ichniowski, C., Shaw, K., 2007. How does information technology affect productivity? plant-level comparisons of product innovation, process improvement, and worker skills. *The Quarterly Journal of Economics* 122, 1721–1758.
- Beaudry, P., Green, D. A., Sand, B. M., 2014. The declining fortunes of the young since 2000. *American Economic Review* 104, 381–86.
- Bena, J., Simintzi, E., 2018. Machines could not compete with chinese labor: Evidence from u.s. firms' innovation. working paper .
- Berman, E., Bound, J., Griliches, Z., 1994. Changes in the Demand for Skilled Labor within U. S. Manufacturing: Evidence from the Annual Survey of Manufactures. *Quarterly Journal of Economics* 109, 367–397.
- Bessen, J., Goos, M., Salomons, A., van den Berge, W., 2019. Automatic reaction – What happens to workers at firms that automate? working paper .
- Bond, P., Edmans, A., Goldstein, I., 2012. The real effects of financial markets. *Annual Review of Financial Economics* 4, 339–360.
- Bound, J., Jaeger, D., 2000. Do compulsory attendance laws alone explain the association between earnings and quarter of birth. *Research in Labor Economics* 19, 83–108.
- Bresnahan, T., 1999. Computerisation and wage dispersion: An analytical reinterpretation. *Economic Journal* 256, 390–415.
- Bresnahan, T. F., Brynjolfsson, E., Hitt, L. M., 2002. Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence. *The Quarterly Journal of Economics* 117, 339–376.
- Brown, L., Call, A., Clement, M., Sharp, N., 2015. Inside the “black box” of sell-side financial analysts. *Journal of Accounting Research* 53, 1–47.
- Brynjolfsson, E., McAfee, A., 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton.
- Brynjolfsson, E., Mitchell, T., Rock, D., 2018. What can machines learn, and what does it mean for occupations and the economy? *AEA Papers and Proceedings* 108, 43–47.
- Castellanos, S., 2020. No ‘ai winter’ in the forecast, tech pioneers say. *Wall Street Journal* February 11.
- Castex, G., Dechter, E., 2014. The changing roles of education and ability in wage determination. *Journal of Labor Economics* 32, 685–710.
- Clement, M., Tse, S., 2005. Financial analyst characteristics and herding behavior in forecasting. *The Journal of Finance* 60, 307–341.

- Cohen, L., Frazzini, A., Malloy, C., 2010. Sell-side school ties. *Journal of Finance* 65, 1409–1437.
- Coleman, B., Merkley, K. J., Pacelli, J., 2020. Man versus machine: A comparison of robo-analyst and traditional research analyst investment recommendations. working paper .
- Conley, T. G., Hansen, C. B., Rossi, P. E., 2012. Plausibly exogenous. *Review of Economics and Statistics* 94, 260–272.
- Cowgill, B., Dell’Acqua, F., Deng, S., Hsu, D., Verma, N., Chaintreau, A., 2020. Biased programmers? or biased data? a field experiment in operationalizing ai ethics. *Proceedings of the 21st ACM Conference on Economics and Computation* .
- Da, Z., Huang, X., 2020. Harnessing the wisdom of crowds. *Management Science* forthcoming.
- Da, Z., Jin, L., Huang, X., 2020. Extrapolative beliefs in the cross-section: What can we learn from the crowds. *Journal of Financial Economics* .
- D’Acunto, F., Prabhala, N., Rossi, A. G., 2019. The promises and pitfalls of robo-advising. *Review of Financial Studies* 32, 1983–2020.
- Dauth, W., Findeisen, S., Suedekum, J., Woessner, N., 2018. Adjusting to robots: Worker-level evidence. *FRB of Minneapolis Working Paper* .
- Deming, D., Kahn, L. B., 2018. Skill requirements across firms and labor markets: Evidence from job postings for professionals. *Journal of Labor Economics* 36, S337–S369.
- Deming, D. J., 2017. The Growing Importance of Social Skills in the Labor Market. *Quarterly Journal of Economics* 132, 1593–1640.
- Deming, D. J., Weidmann, B., 2020. Team players: How social skills improve group performance. working paper .
- Detweiler, G., 2019. 20 financial bloggers share their secrets to running a successful blog. *Forbes* November 2.
- Diether, K. B., Malloy, C. J., Scherbina, A., 2002. Differences of opinion and the cross section of stock returns. *Journal of Finance* 57, 2113–2141.
- Efron, B., Hastie, T., Johnstone, I., Tibshirani, R., 2004. Least angle regression. *The Annals of Statistics* 32, 407–499.
- Emery, D. R., Li, X., 2009. Are the Wall Street analyst rankings popularity contests? *Journal of Financial and Quantitative Analysis* pp. 411–437.
- Erel, I., Stern, L. H., Tan, C., Weisbach, M. S., 2018. Selecting directors using machine learning. Working Paper 24435, National Bureau of Economic Research.

- Fang, L., Yasuda, A., 2009. The Effectiveness of Reputation as a Disciplinary Mechanism in Sell-Side Research. *The Review of Financial Studies* 22, 3735–3777.
- Felten, E. W., Raj, M., Seamans, R., 2018. A method to link advances in artificial intelligence to occupational abilities. *AEA Papers and Proceedings* 108, 54–57.
- Felten, E. W., Raj, M., Seamans, R., 2019. The occupational impact of artificial intelligence: Labor, skills, and polarization, unpublished working paper.
- Frank, M. R., Autor, D., Bessen, J. E., Brynjolfsson, E., Cebrian, M., Deming, D. J., Feldman, M., Groh, M., Lobo, J., Moro, E., Wang, D., Youn, H., Rahwan, I., 2019. Toward understanding the impact of artificial intelligence on labor. *Proceedings of the National Academy of Sciences* 116, 6531–6539.
- Frey, C. B., Osborne, M. A., 2017. The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change* 114, 254–280.
- Goldfarb, A., Taska, B., Teodoridis, F., 2020. Could machine learning be a general purpose technology? Evidence from online job postings. working paper .
- Graetz, G., Michaels, G., 2018. Robots at work. *The Review of Economics and Statistics* 100, 753–768.
- Graham, J. R., Grennan, J., Harvey, C. R., Rajgopal, S., 2018. Corporate culture: The interview evidence, unpublished working paper.
- Graham, J. R., Grennan, J., Harvey, C. R., Rajgopal, S., 2019. Corporate culture: Evidence from the field, unpublished working paper.
- Grennan, J., 2019. A corporate culture channel: How increased shareholder governance decreases firm value, unpublished working paper.
- Grennan, J., Michaely, R., 2020. FinTechs and the market for financial analysis. *Journal of Financial and Quantitative Analysis* forthcoming.
- Groysberg, B., Healy, P. M., Maber, D. A., 2011. What drives sell-side analyst compensation at high-status investment banks? *Journal of Accounting Research* 49, 969–1000.
- Hacamo, I., Kleiner, K., 2020. Competing for talent: Firms, managers, and social networks, unpublished working paper.
- Heckman, J. J., Kautz, T., 2012. Hard evidence on soft skills. *Labour Economics* 19, 451–464.
- Hoffmann, M., Kahn, L., Li, D., 2018. Discretion in hiring. *Quarterly Journal of Economics* 133, 765–800.
- Hong, H., Kacperczyk, M., 2010. Competition and bias. *Quarterly Journal of Economics* 125, 1683–1725.

- Hong, H., Kubik, J. D., 2003. Analyzing the analysts: Career concerns and biased earnings forecasts. *The Journal of Finance* 58, 313–351.
- Humlum, A., 2020. Robot adoption and labor market dynamics. working paper .
- Humlum, A., Meyer, B. B., 2020. Artificial intelligence and college majors. working paper .
- Imbens, G. W., Rubin, D. B., 2015. *Causal Inference in Statistics, Social, and Biomedical Sciences*. New York, NY: Cambridge University Press.
- Kadan, O., Madureira, L., Wang, R., Zach, T., 2012. Analysts' industry expertise. *Journal of Accounting and Economics* 54, 95–120.
- Kambourov, G., Manovskii, I., 2009. Occupational Mobility and Wage Inequality. *Review of Economic Studies* 76, 731–759.
- Katz, L. F., Murphy, K. M., 1992. Changes in relative wages, 1963-1987: Supply and demand factors. *The Quarterly Journal of Economics* 107, 35–78.
- Kempf, E., 2020. The job rating game: Revolving doors and analyst incentives. *Journal of Financial Economics* 135, 41–67.
- Kippersluis, H. V., Rietveld, C. A., 2018. Beyond plausibly exogenous. *Journal of Econometrics* 21, 316–331.
- Konnikova, M., 2014. How headlines change the way we think. *New Yorker* December 17.
- Kratsios, M., Parker, L. E., 2020. America artificial intelligence initiative: Year one annual report. Tech. rep., The White House Office of Science and Technology Policy.
- Lamont, O. A., 2002. Macroeconomic forecasts and microeconomic forecasters. *Journal of Economic Behavior & Organization* 48, 265–280.
- Lazear, E. P., Shaw, K. L., 2007. Personnel economics: The economist's view of human resources. *Journal of Economic Perspectives* 21, 91–114.
- Lee, Y. S., Webb, M., Larsen, B., Cuéllar, M.-F., 2019. How would ai regulation change firms' behavior? Evidence from thousands of managers. working paper .
- Li, D., Raymond, L., Bergman, P., 2020. Hiring as exploration. working paper .
- Ljungqvist, A., Malloy, C., Marston, F., 2009. Rewriting history. *Journal of Finance* 64, 1935–1960.
- Ljungqvist, A., Marston, F., Starks, L. T., Wei, K. D., Yan, H., 2007. Conflicts of interest in sell-side research and the moderating role of institutional investors. *Journal of Financial Economics* 85, 420–456.

- Malmendier, U., Shanthikumar, D., 2014. Do security analysts speak in two tongues? *Review of Financial Studies* 27, 1287–1322.
- Merkley, K., Michaely, R., Pacelli, J., 2017. Does the scope of the sell-side analyst industry matter? An examination of bias, accuracy, and information content of analyst reports. *Journal of Finance* 72, 1285–1334.
- Merkley, K., Michaely, R., Pacelli, J., 2020. Cultural diversity on wall street: Evidence from consensus earnings forecasts. *Journal of Accounting and Economics* .
- Michaels, G., Natraj, A., Van Reenen, J., 2014. Has ICT polarized skill demand? Evidence from eleven countries over twenty-five years. *Review of Economics and Statistics* 96, 60–77.
- Michaely, R., Rubin, A., Segal, D., Vedrashko, A., 2018. Lured by the consensus, unpublished working paper.
- Michaely, R., Womack, K. L., 1999. Conflict of interest and the credibility of underwriter analyst recommendations. *Review of Financial Studies* 12, 653–686.
- Mikhail, M. B., Walther, B. R., Willis, R. H., 1999. Does forecast accuracy matter to security analysts? *The Accounting Review* 74, 185–200.
- Mishra, S., 2020. The ai index 2019 annual report. Tech. rep., AI Index Steering Committee, Human-Centered AI Institute, Stanford University.
- Robinson, P., 2019. Writing a good headline isn't as simple as it sounds. *ACES: The Society for Editing* August 12.
- Webb, M., 2020. The impact of artificial intelligence on the labor market, unpublished working paper.
- Zator, M., 2020. Digitization and automation: Firm investment and labor outcomes. working paper .

Figure 1.

AI Intensity and Career Decisions of Analysts

This figure plots the relationship between AI data intensity and analysts' decisions to switch careers at the analyst-quarter level. This figure on the left focuses on the quitting decision of any analyst and the figure on the right the quitting decision of analysts in the top 25th percentile of accuracy prior to quitting. Each dot shows the average quitting decision for a given percent of the stocks that they cover having high AI intensity, after controlling for the workload (total number of stocks covered), experience, reputation (times nominated as an All-star by Institutional Investor magazine), and time fixed effects. The plotted line represents the best linear approximation to the conditional expectation function.




Figure 2.

Analyst Quit Rates and Career Transitions Over Time with Controls

This figure plots how analysts' decisions to switch careers changes over time after controlling for additional factors. The top plots focus on quit rates and the bottom plots focus on the next job that analysts take. Each dot shows the average quitting decision for a given time period, after controlling for the workload (total number of stocks covered), experience, reputation (times nominated as an All-star by Institutional Investor magazine), and employer fixed effects. The plotted line represents the best linear approximation to the conditional expectation function.




Table I

Summary Statistics

This table presents summary statistics for the main dependent and independent variables. The main sample period is limited to 2010Q1 to 2016Q4.

	Freq.	Mean	Median	Std. Dev.	Obs.
	(1)	(2)	(3)	(4)	(5)
Analyst-level variables					
Quits	Q	3.9%	0.0%	19.3%	73,778
Quits & among top 25% for accuracy	Q	0.7%	0.0%	8.5%	73,778
Quits & Institutional Investor All-star	Q	0.1%	0.0%	2.9%	73,778
Stocks covered with high AI intensity	Q	23%	13%	29%	73,778
Stocks covered	Q	8.6	8.0	6.2	73,778
Industries covered	Q	2.5	2.0	1.8	73,778
General experience (years)	Q	5.7	6.0	3.6	73,778
Times nominated as Institutional Investor All-star	Q	0.16	0.00	0.90	73,778
Stock-level variables					
Questions on earnings call	Q	3.59	3.28	1.81	55,260
Question complexity on earnings call	Q	55.2	53.5	17.6	55,260
Easy-to-measure question topic	Q	0.85	1.00	0.35	55,260
Hard-to-measure question topic	Q	0.19	0.00	0.39	55,260
Analyst meetings	Q	11.5	5.0	15.7	72,149
Analyst meetings with management	Q	9.5	4.0	13.7	72,149
Analyst meetings with institutional investors	Q	2.0	0.0	3.5	72,149
Analyst-stock-level variables					
Initiates stock coverage	Q	28.1%	0.0%	45.0%	245,284
Ends stock coverage	Q	4.76%	0.00%	21.3%	633,644
Accuracy of earnings forecast (as % of price)	Q	0.013	0.005	0.024	633,644
Bias of earnings forecast (as % of price)	Q	0.007	0.021	0.002	633,644
Accuracy of earnings forecast (as % of consensus)	Q	0.70	0.35	0.97	601,835
Bias of earnings forecast (as % of consensus)	Q	0.45	0.14	1.18	601,835
Boldness of earnings forecast	Q	0.31	0.29	0.23	601,835
AI intensity	Q	22.6	10.0	36.4	633,644
Key news events	Q	19.2	14.0	27.1	633,644
Total analysts covering stock	Q	12.6	11.4	7.3	633,644
Firm size	Q	15.1	15.1	1.66	633,644
Daily return volatility	Q	34.2%	29.6%	17.7%	633,644
Mean monthly return	Q	1.3%	1.1%	5.9%	633,644
Amihud's illiquidity ratio	Q	5.4	36.1	275.0	633,644
Log market-to-book ratio	Q	0.91	0.84	0.45	633,644
Volatility of ROE	Q	52.4%	0.2%	322.1%	633,644
Profitability	Q	2.7%	2.8%	3.6%	633,644
Member of S&P 500	Q	35.2%	0.0%	47.8%	633,644

Table II**Career Choices of Analysts**

This table presents OLS and IV estimates of the relationship between AI data intensity and analysts' decisions to switch careers at the analyst-quarter level. In Columns (1)–(2), the dependent variable is an indicator variable for if an analyst quits being a sell-side analyst. In Columns (3)–(4), the dependent variable is an indicator for if an analyst in the top 25th percentile of accuracy in the time period prior to quitting. In Columns (5)–(6), the dependent variable is an indicator for if an analyst nominated by Institutional Investor Magazine as an All-star quits the profession. The focal independent variable is the percent of stocks covered by the analyst with high AI intensity, defined as being in the upper quartile of social media posts. The IV is the percent of stocks covered by the analyst with below median headline length in the USA Today in that quarter. Additional control variables include analyst work experience (early career, general experience), analyst reputation (times nominated as an All-star), and analyst workload (stocks covered, industries covered). Below the coefficient estimates are robust standard errors clustered at the analyst level. The fixed effects are used in both the OLS and IV regressions and are denoted at the bottom of the table. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Analyst quits		Accurate analyst quits		All-star analyst quits	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. OLS						
Stocks covered with high AI intensity	0.015** (0.006)	0.009 (0.006)	0.024*** (0.006)	0.012** (0.006)	0.006 (0.005)	-0.001 (0.005)
Adjusted R^2	3.2%	5.4%	0.6%	1.4%	0.2%	0.5%
	Analyst quits		Accurate analyst quits		All-star analyst quits	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel B. IV						
Stocks covered with high AI intensity	0.029** (0.014)	0.029* (0.016)	0.054*** (0.017)	0.041** (0.019)	0.002 (0.008)	-0.007 (0.011)
First-stage F statistic	2135.9	1528.4	2135.9	1528.4	2135.9	1528.4
t -statistic on IV	46.2	39.1	46.2	39.1	46.2	39.1
Analyst-quarter observations	73,778	73,735	73,778	73,735	73,778	73,735
Unique analysts	5,938	5,909	5,938	5,909	5,938	5,909
Additional analyst controls	Y	Y	Y	Y	Y	Y
Main industry covered fixed effects	Y	Y	Y	Y	Y	Y
Year fixed effects	Y	Y	Y	Y	Y	Y
Employer fixed effects	N	Y	N	Y	N	Y

Table III**Robustness: Career Choices of Analysts**

This table presents IV estimates of the relationship between AI data intensity and analysts' decisions to switch careers at the analyst-quarter level. In Columns (1)–(2), the dependent variable is an indicator variable for if an analyst quits being a sell-side analyst. In Columns (3)–(4), the dependent variable is an indicator for if an analyst in the top 25th percentile of accuracy in the time period prior to quitting. In Columns (5)–(6), the dependent variable is an indicator for if an analyst nominated by Institutional Investor Magazine as an All-star quits the profession. The primary independent variable of interest is the percent of stocks covered by the analyst with high AI intensity, defined as being in the upper quartile of AI-intensity. The IV is the percent of stocks covered by the analyst with below median headline length in the USA Today in that quarter. Additional control variables include analyst work experience, analyst reputation, and analyst workload. Below the coefficient estimates are robust standard errors clustered at the analyst level. The fixed effects are used in all IV regressions and are denoted at the bottom of the table. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Analyst quits		Accurate analyst quits		All-star analyst quits	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Remove Tail of Stocks Covered						
Stocks covered with high AI intensity	0.034***	0.037***	0.053***	0.048***	0.015	0.008
	(0.012)	(0.014)	(0.015)	(0.018)	(0.011)	(0.015)
First-stage F statistic	1792.3	1210.1	1792.3	1210.1	1792.3	1210.1
t -statistic on IV	42.3	34.8	42.3	34.8	42.3	34.8
Analyst-quarter observations	57,252	57,218	57,252	57,218	57,252	57,218
Panel B. IQR of High AI Intensity						
Stocks covered with high AI intensity	0.052***	0.052**	0.061***	0.058**	0.002	-0.007
	(0.018)	(0.021)	(0.022)	(0.025)	(0.008)	(0.011)
First-stage F statistic	1111.1	857.8	1111.1	857.8	1111.1	857.8
t -statistic on IV	33.3	29.3	33.3	29.3	33.3	29.3
Analyst-quarter observations	39,959	39,922	39,959	39,922	39,959	39,922
Panel C. Alt. Definition of AI Intensity						
Stocks covered with high AI intensity	0.035**	0.033**	0.064***	0.046**	0.003	-0.008
	(0.017)	(0.018)	(0.020)	(0.021)	(0.010)	(0.013)
First-stage F statistic	567.7	549.3	567.7	549.3	567.7	549.3
t -statistic on IV	23.8	23.4	23.8	23.4	23.8	23.4
Analyst-quarter observations	73,778	73,735	73,778	73,735	73,778	73,735
Panel D. Alt. Accuracy Threshold						
	Top 15% accuracy		Top 20% accuracy		Top 33% accuracy	
	(1)	(2)	(3)	(4)	(5)	(6)
Stocks covered with high AI intensity	0.057***	0.051**	0.056***	0.045**	0.063***	0.051**
	(0.018)	(0.021)	(0.017)	(0.020)	(0.017)	(0.020)
First-stage F statistic	2135.9	1528.4	2135.9	1528.4	2135.9	1528.4
t -statistic on IV	46.2	39.1	46.2	39.1	46.2	39.1
Analyst-quarter observations	73,778	73,735	73,778	73,735	73,778	73,735
Additional analyst controls	Y	Y	Y	Y	Y	Y
Main industry covered fixed effects	Y	Y	Y	Y	Y	Y
Year fixed effects	Y	Y	Y	Y	Y	Y
Employer fixed effects	N	Y	N	Y	N	Y

Table IV

Next Job After Being an Analyst

This table presents IV estimates of the relationship between AI intensity and the next job the analyst takes after leaving the profession. The dependent variables are a research job, a non-research job, a FinTech job, and a new position at the same employer, respectively. The primary independent variable of interest is the percent of stocks covered by the analyst with high AI intensity, defined as being in the upper quartile of AI-intensity. The IV is the percent of stocks covered by the analyst with below median headline length in the USA Today in that quarter. Additional control variables include experience (years at firm), reputation (times nominated as an All-star analyst), and workload (number of stocks covered, number of industries covered). Each regression includes year fixed effects. Below the coefficient estimates are robust standard errors clustered at the analyst level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Research job (1)	Non-research job (2)	FinTech job (3)
Stocks covered with high AI intensity	-0.134* (0.079)	0.182** (0.086)	-0.016 (0.046)
First-stage F statistic	47.5	47.5	47.5
t -statistic on IV	6.9	6.9	6.9
Analyst observations	209	209	209
Additional analyst controls	Y	Y	Y
Year fixed effects	Y	Y	Y

Table V**Stock Coverage by Analysts**

This table presents IV estimates of the relationship between AI-analyst intensity and actual sell-side analysts' decisions to change the stock they cover at the analyst-stock-quarter level. In Panel A, the dependent variable is an indicator variable for if an analyst initiates coverage of stock in the next quarter. In Panel B, the dependent variable is an indicator variable for if an analyst quits covering a stock in that quarter. The primary independent variable of interest, AI intensity, which measures the quantity of social media posts analyzed to generate AI-powered stock picks in quarter t for stock i . The IV for AI intensity is an indicator variable for whether the newspaper headlines in the USA Today for that stock were shorter than the median headline length in that quarter. Additional analyst controls include analyst experience, workload (stocks covered, industries covered), and reputation. Additional stock controls include key news events, total analyst coverage, firm size, average monthly returns, volatility, market-to-book ratio, volatility of ROE, Amihud's illiquidity ratio, profitability, an indicator for being in the S&P 500, and counts of key newspaper headline words. Below the coefficient estimates are robust standard errors clustered at the stock level. The fixed effects used in both the OLS and IV regressions are denoted at the bottom of the table. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Initiate coverage		Stop coverage	
<i>Panel A. OLS</i>	(1)	(2)	(3)	(4)
AI intensity	-0.042***	-0.057***	-0.004***	-0.013***
	(0.007)	(0.008)	(0.002)	(0.003)
Adjusted R^2	20.7%	15.2%	15.8%	2.9%
	Initiate coverage		Stop coverage	
<i>Panel B. IV</i>	(1)	(2)	(3)	(4)
AI intensity	-0.043***	-0.066***	0.018***	0.090***
	(0.013)	(0.025)	(0.006)	(0.029)
Additional analyst & stock controls	Y	Y	Y	Y
Main industry covered fixed effects	Y	Y	Y	Y
Year fixed effects	Y	Y	Y	Y
Analyst fixed effects	Y	N	Y	N
Stock fixed effects	N	Y	N	Y
First-stage F statistic	76.8	54.0	108.7	72.4
t -statistic on IV	8.8	7.4	10.4	8.5
Analyst-stock-quarter obs.	244,210	244,975	632,962	633,225

Table VI

Robustness: Stock Coverage by Analysts

This table summarizes tests related to potential violations of the exclusion restriction assumption. Panel A examines the direct effect of the IV on stock coverage decisions for the zero-first-stage and remaining group. The zero-first-stage is defined as stocks in the lowest decile of market capitalization. Panel B reports the IV and plausibly exogenous IV estimates. The plausibly exogenous IV is estimated using Conley et al. (2012) and the prior distribution with Imbens and Rubin uncertainty follows the procedure in Kippersluis and Rietveld (2018). The focal independent variable is AI intensity, which measures the quantity of social media posts analyzed to generate AI-powered stock picks in quarter t for stock i . The IV for AI intensity is an indicator variable for whether the newspaper headlines in the USA Today for that stock were shorter than the median headline length in that quarter. Additional analyst controls include analyst experience, workload (stocks covered, industries covered), and reputation. Additional stock controls include key news events, total analyst coverage, firm size, average monthly returns, volatility, market-to-book ratio, volatility of ROE, Amihud's illiquidity ratio, profitability, an indicator for being in the S&P 500, and counts of key newspaper headline words. Standard errors clustered by stock are reported below the coefficient estimates. The fixed effects used in both Panel A and B are denoted at the bottom of the table. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable =			
	Initiate coverage (1)	Stop coverage (2)	Stop coverage (3)	Stop coverage (4)
<i>Panel A. The effect of IV on stock coverage</i>				
Zero-first stage group	0.024 (0.187)	0.191 (0.241)	-0.003 (0.076)	0.083 (0.099)
Observations	23,736	24,265	62,887	63,507
Remaining group	-0.076*** (0.021)	-0.067*** (0.025)	0.024*** (0.008)	0.041*** (0.012)
Observations	219,706	220,607	569,654	569,694
<i>Panel B. The effect of AI intensity on stock coverage</i>				
IV	-0.043*** (0.013)	-0.066*** (0.025)	0.018*** (0.006)	0.090*** (0.029)
Plausibly Exogenous IV, $\zeta \sim (0.0, \Omega_\zeta)$	-0.043** (0.019)	-0.066* (0.039)	0.018** (0.009)	0.090** (0.040)
Additional analyst & stock controls	Y	Y	Y	Y
Main industry covered fixed effects	Y	Y	Y	Y
Year fixed effects	Y	Y	Y	Y
Analyst fixed effects	Y	N	Y	N
Stock fixed effects	N	Y	N	Y
Observations	244,210	244,975	632,962	633,225

Table VII**Analysts' Efforts During Earnings Conference Calls**

This table presents IV estimates of the relationship between AI intensity and analysts' actions during earnings conference calls at the stock-quarter level. Panel A reports coefficient estimates about analyst questions. In Columns (1)–(2) of Panel A, the dependent variable is total analyst questions. In Columns (3)–(4) of Panel A, the dependent variable is question complexity. Panel B reports coefficient estimates related to the content of analyst questions. In Columns (1)–(2) of Panel B, the dependent variable is an indicator for a question about an easy-to-measure topic for AI and in Columns (3)–(4) of Panel B, the dependent variable is an indicator for a question about a harder-to-measure topic for AI. The primary independent variable of interest, AI intensity, which measures the quantity of social media posts analyzed to generate AI-powered stock picks in quarter t for stock i . The IV for AI data intensity is an indicator variable for whether the newspaper headlines in the USA Today for that stock were shorter than the median headline length in that quarter. Additional controls include: key events for that stock, total analyst coverage, firm size, average monthly returns, volatility, market-to-book ratio, volatility of ROE, Amihud's illiquidity ratio, profitability, an indicator for being in the S&P 500, and counts of key newspaper headline words. Below the coefficient estimates are robust standard errors clustered at the stock level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Total questions		Question complexity	
<i>Panel A. Analyst questions</i>	(1)	(2)	(3)	(4)
AI intensity	0.082** (0.042)	0.077* (0.044)	0.088** (0.043)	0.098** (0.044)
	Easy-to-measure topic		Hard-to-measure topic	
<i>Panel B. Content of analyst questions</i>	(1)	(2)	(3)	(4)
AI intensity	-0.094* (0.051)	-0.124*** (0.029)	0.104* (0.060)	0.041 (0.032)
Additional stock controls	Y	Y	Y	Y
Year fixed effects	Y	Y	Y	Y
Industry fixed effects	Y	Y	Y	Y
Stock fixed effects	Y	N	Y	N
First-stage F statistic	131.5	207.1	131.5	207.1
t -statistic on IV	11.5	14.4	1.5	14.4
Stock-quarter observations	55,057	55,260	55,057	55,260

Table VIII**Analysts' Meetings with Management and Investors**

This table presents IV estimates of the relationship between analysts' time spent doing work that relies on their soft skills and the exposure to high AI intensity stocks in their portfolio. In Columns (1)–(2), the dependent variable is the number of meeting an analyst participates in. In Columns (3)–(4), the dependent variable is the number of meetings with institutional investors. In Columns (5)–(6), the dependent variable is the number of meetings with management. The focal independent variable is the percent of stocks covered by the analyst with high AI intensity, defined as being in the upper quartile of AI intensity. The IV is the percent of stocks covered by the analyst with below median headline length in the USA Today in that quarter. Additional control variables include analyst's work experience, workload (number of stocks covered, number of industries covered), and reputation. Below the coefficient estimates are robust standard errors clustered at the analyst level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Total analyst meetings		Meetings with institutional investors		Meetings with management	
	(1)	(2)	(3)	(4)	(5)	(6)
Stocks covered with high AI intensity	0.166*** (0.012)	0.137*** (0.013)	0.159*** (0.013)	0.139*** (0.015)	0.150*** (0.012)	0.121*** (0.013)
Additional analyst controls	Y	Y	Y	Y	Y	Y
Year fixed effects	Y	Y	Y	Y	Y	Y
Main industry covered fixed effects	Y	Y	Y	Y	Y	Y
Employer fixed effects	N	Y	N	Y	N	Y
First-stage <i>F</i> statistic	2402.6	1699.0	2402.6	1699.0	2402.6	1699.0
<i>t</i> -statistic on IV	49.0	41.2	49.0	41.2	49.0	41.2
Analyst-quarter observations	72,149	72,112	72,149	72,112	72,149	72,112

Table IX**Quality of Analysts' Reports: Analyst-Stock Pairs**

This table presents IV estimates of the relationship between AI intensity and the quality of analysts' reports at the analyst-stock-quarter level. In Columns (1)–(2), the dependent variables are analyst accuracy and bias as a percent of the consensus forecast among all analysts tracking stock i in quarter t . In Columns (3)–(4), the dependent variables are analyst accuracy and bias as a percent of the stock price at the close of the previous quarter. In Column (5), the dependent variable is the boldness of the analyst's forecast. The focal independent variable is AI intensity which measures the quantity of social media posts analyzed to generate AI-powered stock picks for stock i in quarter t . The IV for AI intensity is an indicator variable for whether the newspaper headlines in the USA Today for that stock were shorter than the median headline length in that quarter. Additional analyst controls include total number of stocks covered, total number of industries covered, general experience as an analyst, and reputation. Additional stock controls include the total key news events, total analyst coverage, firm size, average monthly returns, volatility, market-to-book ratio, volatility of ROE, Amihud's illiquidity ratio, profitability, an indicator for being in the S&P 500, and counts of key newspaper headline words. Below the coefficient estimates are robust standard errors clustered at the stock level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable =				
	As % of consensus		As % of stock price		Bold
	Accuracy	Bias	Accuracy	Bias	forecast
	(1)	(2)	(3)	(4)	(5)
AI intensity	0.141** (0.072)	-0.102 (0.065)	0.065* (0.039)	-0.123** (0.050)	0.142** (0.064)
Additional analyst & stock controls	Y	Y	Y	Y	Y
Year fixed effects	Y	Y	Y	Y	Y
Main industry covered fixed effects	Y	Y	Y	Y	Y
Analyst-by-stock fixed effects	Y	Y	Y	Y	Y
Employer fixed effects	Y	Y	Y	Y	Y
First-stage F statistic	72.6	72.6	73.2	73.2	73.1
t -statistic on IV	8.5	8.5	8.6	8.6	8.6
Analyst-stock-quarter observations	594,129	594,129	625,639	625,639	597,281

Table X

Quality of Analysts' Reports in Aggregate

This table presents IV estimates of the relationship between AI intensity and the quality of analysts' reports at the stock-quarter level. In Panel A, B, and C, the dependent variables are accuracy, bias, and the consensus forecast error. Analysts' forecasts are aggregated to the stock level by taking the mean in Columns (1) and (3) or by taking the median in Columns (2) and (4). The focal independent variable is AI intensity, which measures the quantity of social media posts analyzed to generate AI-powered stock picks in quarter t for stock i . The IV for AI intensity is an indicator variable for whether the newspaper headlines in the USA Today for that stock were shorter than the median headline length in that quarter. Additional controls include key events for that stock, total analyst coverage, firm size, average monthly returns, volatility, market-to-book ratio, volatility of ROE, Amihud's illiquidity ratio, profitability, an indicator for being in the S&P 500, and counts of key newspaper headline words. Below the coefficient estimates are robust standard errors clustered at the stock level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	As % of consensus		As % of stock price	
	Mean	Median	Mean	Median
<i>Panel A. Dependent variable = Accuracy</i>				
AI intensity	0.077 (0.063)	0.064 (0.057)	0.053** (0.027)	0.070*** (0.024)
<i>Panel B. Dependent variable = Bias</i>				
AI intensity	-0.067 (0.057)	-0.033 (0.051)	-0.064* (0.036)	-0.083** (0.035)
<i>Panel C. Dependent variable = Consensus Forecast Error</i>				
AI intensity	-0.108*** (0.032)			
Additional stock controls	Y	Y	Y	Y
Stock fixed effects	Y	Y	Y	Y
Year fixed effects	Y	Y	Y	Y
First-stage F statistic	104.1	104.1	111.3	111.3
t -statistic on IV	10.2	10.2	10.6	10.6
Stock-quarter observations	83,237	83,237	91,095	91,095

Table XI**Analyst Compensation: Evidence from Market Reactions to Recommendations**

This table presents estimates of abnormal returns and volume following analysts' recommendation revisions, where revisions are limited to an upgrade to a buy or strong buy or a downgrade to a sell or strong sell. Returns are in excess of benchmark portfolios matched on size, book-to-market, and momentum. Downgrades have been multiplied by -1 to reflect the opposite predicted direction of stock returns. Log volume is relative to a market model. Downgrades are not multiplied by -1 for volume. The primary independent variable of interest, AI intensity, which measures the quantity of social media posts analyzed to generate AI-powered stock picks in quarter t for stock i . ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dep. var. = Excess Returns		Dep. var. = Excess Volume	
	[0,1] (1)	[0,5] (2)	[0,1] (3)	[0,5] (4)
Reaction to Analyst Recommendations				
AI intensity	-0.24%*** (0.05%)	-0.27%*** (0.05%)	-0.47*** (0.10)	-0.034* (0.19)
Time fixed effects	Y	Y	Y	Y
Analyst fixed effects	Y	Y	Y	Y
Recommendation observations	39,454	39,454	39,454	39,454

INTERNET APPENDIX

Appendix A. Variable Definitions

We use data from IBES, CRSP, Compustat, Capital IQ, and Thomson Reuters to construct our sell-side analyst sample. Continuous variables are winsorized at the 1st and 99th percentile to mitigate the influence of extreme observations. Definitions are as follows:

Total analysts covering stock is the number of analysts covering stock i in quarter t . ($NUMEST$)

Firm Size is the logarithm of stock i 's market capitalization at the end of quarter t . ($\log(PRCC_F \times CSHO)$)

Daily Return Volatility is the annualized variance of daily raw returns of stock i in quarter t . ($\sigma_{RET} \times \sqrt{252}$).

Mean Monthly Return is the average monthly return on stock i in quarter t . ($R\bar{E}T$)

Log Market-to-book = $\log\left(\frac{PRCC_F \times CSHO + DLC + DLTT + PSTKL - TXDITC}{AT}\right)$

Amihud's Illiquidity Ratio is $ILLIQ = \frac{1}{N} \sum_{t=1}^T \frac{|r_t|}{V_t}$ where T is the number of days depending on the frequency (weekly, quarterly, or annual) and V_t is the dollar-weighted volume of trades and r_t is the returns (?).

Return on equity (ROE) = $\frac{NI}{SEQ_{t-1}}$

Volatility of ROE comes from estimating an AR(1) model for each stock's ROE using a rolling, 10-year series of the company's valid annual ROEs. The variance of the residuals from this regression is the volatility of ROE.

Profitability = $\frac{OIBDP}{AT}$

Member of S&P 500 is an indicator variable that takes the value of one if stock i is included in the S&P 500 index in quarter t .

Key news events come from Capital IQ's database, which provides summaries of material news and events that may affect the market value of securities. It monitors over 100 key development types including executive changes, M&A rumors, SEC inquiries, etc... Each key development item includes announcement date and type. For each time period, we create a count of the total number of key developments reported. We also create counts for the total number of key developments in the following categories: payout events (*keydeveventtypeid* equal to 45, 46, 47, 94, 151, 152, 213, 214, 230, 231, 232, 233 and 234), credit events (*keydeveventtypeid* equal to 7, 68, 69, and 89), activist interventions (*keydeveventtypeid* equal to 156, 157, 158, 159, 161, 166, 169, 175, 179, 181, 185, 188, 189, 201, 203, 204), and other events (*keydeveventtypeid* equal 1, 3, 5, 11, 16, 21, 22, 23, 24, 25, 26, 27, 29, 31, 32, 41, 43, 44, 63, 73, 80, 81, 82, 101, 102, 137, 224, 225). The other events category encompasses M&A transactions, labor-related announcements, reorganizations and spin-offs, legal issues, write-offs and impairments, and changes to corporate guidance.

Key newspaper words are the twenty words identified using variable selection techniques as having power to explain headline length.

Stock Experience is the number of years analyst j covered stock i .

General Experience is the number of years since the analyst first appeared in the IBES database.

Number of Stocks Covered is the total number of unique stocks covered by the analyst during the year.

Number of Industries Covered is the total number of unique two-digit SIC industries covered by the analyst during the year.

AI intensity is the number of social media posts discussing stock i in quarter t which we obtain from TipRanks.

Percent of stocks with high AI data intensity is the number of high AI intensity stocks (defined as being in the top quartile of AI intensity in a given quarter) divided by the

total number of stocks covered by analyst j in quarter t .

Analyst Leaves is an indicator variable for if analyst j (defined as a unique *ANALYS* in IBES pneumonics) stops appearing in the IBES dataset altogether. Given that our analyst data extend beyond the sample period for our AI data, we can calculate the number of analysts who quit even in the final quarter. This ensures that truncation errors are not driving our results.

Analyst Leaves and Is in Top 25% of Accuracy is an indicator variable for if analyst j stops appearing in the IBES dataset altogether and was in the top 25% of accuracy across all analysts in quarter $t - 1$. Accuracy is based on the absolute value of the forecast error relative to the stock price in the previous quarter. We take the average accuracy across the stocks covered by that analyst in that quarter.

Analyst Leaves and Is Institutional Investor All-star is an indicator variable for if analyst j stops appearing in the IBES dataset altogether and who was previously nominated by Institutional Investor magazine as an All-star analyst.

Times Nominated is the number of times an analyst was nominated as an All-star by Institutional Investor Magazine as of the quarter the analyst's forecast is made.

Research Job is identified from LinkedIn resume data and internet searches. This is an indicator variable for the analyst's next job being as a buy-side analyst or in the asset management industry.

Non-research Job is identified from LinkedIn resume data and internet searches. This is an indicator variable for the analyst's next job being another role in finance, strategy, consulting, investor relations, or at the same employer but a different role.

FinTech Job is identified from LinkedIn resume data and internet searches. This is an indicator variable for the analyst taking a FinTech job.

Analyst Initiates Stock Coverage is an indicator variable for if analyst j (defined as a unique *ANALYS* in IBES pneumonics) begins covering an stock not previously covered

in the IBES database.

Analyst Ends Stock Coverage is an indicator variable for if analyst j (defined as a unique *ANALYS* in IBES pneumonics) ends covering an stock in the IBES database. Given that our analyst data extend beyond the sample period for our AI data, we can calculate the changes in coverage even in the final quarter.

Questions on Earnings Call is the average number of questions asked per analyst on the conference call based on Capital IQ transcript data. To identify participants on a conference call as analysts, we limit the sample to *speakertypeid* = 3 and *transcriptpresentationtypename* = *Final*.

Question Complexity on Earnings Call is the average total word count in the questions asked by analysts based on Capital IQ transcript data. To identify participants on a conference call as analysts, we limit the sample to *speakertypeid* = 3 and *transcriptpresentationtypename* = *Final*.

Easy-to-measure Earnings Topic is an indicator variable for if an analyst's earnings call question contains the word "sale", "margin", "price", or "capital."

Hard-to-measure Earnings Topic is an indicator variable for if an analyst's earnings call question contains the word "adapt," "brand," "engage," or "technology."

Analyst Meetings with Management is a count of the number of meetings analyst and management from the company of a particular stock met in a given quarter. This data comes from Capital IQ and the event is coded as: *keydeventtypeid* = 51 (Company conference presentation)

Analyst Meetings with Institutional Investors is a count of the number of meetings analyst and investors had about a particular stock in a given quarter. This data comes from Capital IQ and the events are coded as: *keydeveventtypeid* = 192 (Analyst/Investor Day) and *keydeventtypeid* = 50 (Shareholder/Analyst Calls)

Analyst Meetings is the sum of analyst meetings with management and meetings with investors.

Accuracy of Forecast is the absolute value of the signed forecast error (i.e., the difference between the analyst's forecast and the actual EPS) divided by either the absolute value of the consensus EPS for a stock i in quarter t or divided by the closing price for stock i in quarter $t - 1$. To match the definition of accuracy used in Hong and Kacperczyk (2010), we use EPS from Compustat rather than IBES. To construct our various measures of accuracy, we use diluted, U.S. currency quarterly earnings per share (EPS) forecasts from one to eight quarters out. We include in our set of forecasts those that are original forecasts, announced confirmations of previous forecasts, and revised forecasts. For ease in interpretation, we multiply by -1 so that bigger numbers represent improvements in accuracy.

Bias of Forecast is the signed forecast error (i.e., the difference between the analyst's forecast and the actual EPS) divided by either the absolute value of the consensus EPS for a stock i in quarter t or divided by the closing price for stock i in quarter $t - 1$. To match the definition of bias used in Hong and Kacperczyk (2010), we use EPS from Compustat rather than IBES. To construct our various measures of bias, we use diluted, U.S. currency quarterly earnings per share (EPS) forecasts from one to eight quarters out. We include in our set of forecasts those that are original forecasts, announced confirmations of previous forecasts, and revised forecasts.

Boldness of Forecast is the percent of forecast revisions for a given quarter t for a stock i that are bold. We follow the construction in Clement and Tse (2005) and define bold as an indicator variable for each analyst j 's forecast revision for stock i in quarter t . It is equal to 1 if analyst j 's forecast is either above or below both the analyst's prior forecast and the mean forecast immediately before the forecast revision, and 0 otherwise.

Consensus Forecast Error is the difference between the mean consensus forecast and the actual EPS for a given quarter t for a stock i divided by the closing price for stock i

in quarter $t - 1$. Following what is standard in the literature, we use EPS from Compustat rather than IBES.

Appendix B. Additional Tables and Figures

Figure B.1.

TipRanks Smart Score stock analysis tool.

This figure shows the summary page for Facebook from TipRanks Smart Score stock analysis tool. The Smart Score provides an indication of whether the stock will outperform or underperform based on eight factors extracted from TipRanks unique datasets. The Smart Score makes use of artificial intelligence (AI) to generate a rating and does not involve human intervention.




Figure B.2.

Analyst Quit Rates Over Time

This figure is a scatter plot of the raw quarterly quit rate for analysts who leave the profession over time. The plotted line represents the simple linear approximation.




Figure B.3.

Analyst Career Transitions Over Time (Detailed Jobs)

This figure plots how analysts' decisions about the next job they take changes over time. Each dot shows the average movement into a new job for a given time period, after controlling for the workload (total number of stocks covered), experience, and reputation (times nominated as an All-star by Institutional Investor magazine). The plotted line represents the best linear approximation to the conditional expectation function. The figure on the top, left is for asset management jobs and the figure on the top, right is for investor relations jobs. The figure on the bottom left is for other finance roles and the figure on the bottom right is for consulting jobs.




Figure B.4.

AI Intensity and Analyst Participation in Events

This figure plots the relationship between AI intensity and analysts' decisions to attend or host meetings with corporate management or institutional investors. Each dot shows the average number of meetings attended for a given percent of the stocks that the analyst covers having high AI intensity, after controlling for the workload (total number of stocks covered and industries covered), experience, reputation (times nominated as an All-star by Institutional Investor magazine), industry (main industry of coverage by analyst), and year. The plotted line represents the best linear approximation to the conditional expectation function.




Table B.1

Example Newspaper Headlines for Wells Fargo

This table presents example headlines for Wells Fargo from September 2016. The articles have similar content, viz., information about Wells Fargo's scandal and settlement with the Consumer Financial Protection Bureau. The table also includes headline length and orders headlines from shortest to longest.

Headline and headline length	
Wells Fargo Is Getting Heat	27
Trust Was Broken at Wells Fargo	31
Wells Fargo Fined for Sales Scam	32
Wild West at Wells Fargo: Our View	34
Wells Fargo Fined \$185m for Fake Accounts	41
Wells Fargo Fined \$185m; 5,300 Were Fired	41
Wells Fargo to Pay \$185 Million Settlement	42
5,300 Wells Fargo Staff Fired Over Bogus Accounts	49
Wells Fargo Fined \$185m for Unauthorized Accounts	49
Wells Fargo to Pay \$185 Million Fine Over Sales Tactics	55
Wells Fargo Fined \$185m for Opening Unauthorized Accounts	57
Wells Fargo to Pay \$185 Million Fine Over Account Openings	58
Wells Fargo Fined \$185m for Fake Accounts; 5,300 Were Fired	59
What Wells Fargo's \$185 Million Settlement May Mean for You	59
Wells Fargo Fined \$185 Million for Improper Account Openings	60
Wells Fargo Fined \$185m for Unauthorized Accounts; Fires 5,300	62
Wells Fargo Fined \$185 Million Over Unwanted Customer Accounts	62
Wells Fargo Fined \$185m for Unauthorized Accts That Hurt Customers	66
Wells Fargo Cuts Bank Sales Goals After \$185m Fine for Fake Accounts	68
Wells Fargo CEO Defends Bank Culture, Lays Blame With Bad Employees	68
How Wells Fargo's High-Pressure Sales Culture Spiraled out of Control	69
Wells Fargo Fined \$185m Over Unauthorized Accounts That Harmed Customers	72
Wells Fargo to Pay \$185 Million Settlement for 'Outrageous' Sales Culture	73
Wells Fargo Fined \$185m for Unauthorized Accounts; Says It Has Fired 5,300	74
Wells Fargo Fires 5,300 People Over Improper Account Openings; Company Fined \$185m	82
Wells Fargo to Pay \$185 Million to Settle Allegations Its Workers Opened Fake Accounts	86
Wells Fargo CEO John Stumpf Puts on a Clinch: How to Weasel out of Real Accountability	86
Wells Fargo Fires 5,300 People for Opening Millions of Phony Accounts; Company Fined \$185m	90
Wells Fargo Fined \$185 Million for Improper Account Openings; 5,300 People Fired in Connection	94
Wells Fargo Settled Over Its Bogus Accounts, but It Still Faces a Fight From Customers and Ex-Employees	103
Wells Fargo Fired 5,300 Workers for Improper Sales Push. the Executive in Charge Is Retiring With \$125 Million.	111

Table B.2

Headline Length and Firm Characteristics

This table presents OLS estimates in which the dependent variable is headline and the explanatory variables are firm and newspaper characteristics. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dep. var. = Headline length (1)
Log Market-to-Book	0.00 (0.00)
Profitability	-0.53 (0.77)
ROE	0.01 (0.00)
Momentum	1.32* (0.79)
Firm Size	-0.02 (0.05)
Newspaper fixed effect	Y
Adjusted R^2	0.1%
Observations	7,538,452

Table B.3

Headline Length for Value Relevant Events

This table presents OLS estimates in which the dependent variable is headline length and the explanatory variables are value relevant events from Capital IQ's key development data set. These regressions use headlines from the USA Today between 2009 and 2016. In Columns (1) and (2), the focal explanatory variable is an indicator for if a value relevant event occurred that day. In Columns (3) and (4), the focus is on earnings events. We define an earnings announcement as the day of and day after to allow time for a headline to publish. In Columns (5)–(8), we examine non-earnings key events such as payout announcements, targeting by activist investors, or credit upgrades/downgrades. In each regression, controls for log market-to-book, profitability, ROE, momentum and firm size are included. Standard errors clustered by stock are reported below the coefficient estimates. Fixed effects are noted in the bottom rows. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Value-relevant events	Dependent variable = Headline length							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Any key event	0.217 (0.290)	-0.081 (0.097)						
Earnings event			-0.679*** (0.233)	-0.756*** (0.206)				
Positive earnings surprise			-0.198 (0.337)	-0.389 (0.270)				
Negative earnings surprise			-0.827 (0.764)	-0.336 (0.543)				
Non-earnings event					0.424 (0.308)	0.133 (0.098)		
Activist event							0.257 (1.141)	1.050 (0.936)
Other key non-earnings events							0.256 (0.332)	0.032 (0.093)
Firm characteristic controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Quarter fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	No	Yes	No	Yes	No	Yes	No	Yes
Adjusted R^2	9%	23%	9%	23%	9%	23%	9%	23%
Observations	431,710	431,240	431,710	431,240	431,710	431,240	431,710	431,240

Table B.4

LASSO Selection of Words Associated with Headline Length

This table presents estimates connecting common words to headline length. The estimates are based on a LASSO regression. This technique helps with the problem of picking out the relevant words from a larger set (i.e., variable selection) by pushing estimates of some coefficients to be exactly zero. The words are listed in the order in which they are selected to be included in the model. Column (1) shows the LASSO adjusted coefficient estimate for the word, and Column (2) displays the cumulative variance explained when that word is included. Given that the variance explained plateaus toward the end, only the first 20 words selected into the model are listed.

Key headline words	Dep. var. = Headline length (1)	R^2 when variable is included (2)
quarterly	24.89	2.25%
available	10.69	7.42%
annual	5.76	7.94%
stories	-6.21	8.11%
market	4.18	8.34%
talk	-14.78	8.41%
events	-8.25	8.71%
financial	10.47	10.78%
agreement	11.25	10.87%
million	8.84	10.96%
morning	5.07	11.08%
mgmt	-4.32	11.26%
billion	6.79	11.31%
investors	6.53	11.59%
capital	6.95	11.62%
sells	-0.28	11.76%
china	5.32	11.89%
week	7.13	12.22%
fund	5.79	12.37%
bank	4.40	12.37%
Additional word controls	Yes	
Firm characteristic controls	Yes	
Observations	7,538,452	

Table B.5

Summary of Analysts' Next Job

This table summarizes the subsequent jobs held by analysts for a random sample of 398 analysts (approximately 25 per year) that stopped being sell-side analysts between 2004 and 2016.

Next Job	Percent	Frequency
Asset management	19.6%	78
Other finance role	16.1%	64
Buy-side research	15.8%	63
Investor relations	9.3%	37
Same firm, different role	9.3%	37
Entrepreneurship	9.1%	36
Consulting	7.0%	28
Corporate strategy	5.3%	21
FinTech	4.8%	19
Other	3.8%	15

Table B.6**Analysts' Meetings (Alternate Fixed Effects)**

This table presents IV estimates of the relationship between AI-analyst intensity and analysts' time spent doing work that relies on their soft skills. In Column (1), the dependent variable is the number of meeting an analyst participates in. In Column (2), the dependent variable is the number of meetings with institutional investors. In Column (3), the dependent variable is the number of meetings with management. The focal independent variable is the percent of stocks covered by the analyst with high AI intensity, defined as being in the upper quartile of AI intensity. The IV is the percent of stocks covered by the analyst with below median headline length in the USA Today in that quarter. Additional control variables include analyst's work experience, workload (number of stocks covered, number of industries covered), and reputation. Below the coefficient estimates are robust standard errors clustered at the analyst level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Total analyst meetings (1)	Meetings with institutional investors (2)	Meetings with management (3)
Stocks covered with high AI intensity	-0.019 (0.024)	0.106*** (0.029)	-0.049** (0.024)
Additional analyst controls	Y	Y	Y
Year fixed effects	Y	Y	Y
Main industry covered fixed effects	Y	Y	Y
Employer fixed effects	N	N	N
Analyst fixed effects	Y	Y	Y
First-stage F statistic	369.3	369.3	369.3
t -statistic on IV	19.2	19.2	19.2
Analyst-quarter observations	71,648	71,648	71,648

Table B.7

Quality of Analysts' Reports: Analyst-Stock Pairs Details

This table presents IV estimates of the relationship between AI intensity and the quality of analysts' reports at the analyst-stock-quarter level. In Columns (1)–(2), the dependent variables are analyst accuracy and bias as a percent of the consensus forecast among all analysts tracking stock i in quarter t . The IV for AI intensity is an indicator variable for whether the newspaper headlines in the USA Today for that stock were shorter than the median headline length in that quarter. Additional headline variables are counts of key newspaper headline words. Below the coefficient estimates are robust standard errors clustered at the stock level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable =	
	Accuracy	Bias
	(1)	(2)
AI intensity	0.140*	-0.120**
	(0.072)	(0.058)
Key news events	0.010	0.002
	(0.011)	(0.010)
Illiquidity	0.072***	-0.030**
	(0.014)	(0.013)
Volatility	-0.070***	0.089***
	(0.015)	(0.013)
Returns	-0.016***	-0.020***
	(0.005)	(0.004)
Market-to-book	0.083**	-0.001
	(0.037)	(0.002)
Firm Size	0.527***	-0.298***
	(0.081)	(0.070)
Volatility of ROE	0.006	0.016
	(0.021)	(0.015)
Profitability	0.173***	-0.278***
	(0.020)	(0.016)
Analyst coverage	-0.035	0.178***
	(0.032)	(0.020)
Member S&P 500	-0.034	0.055
	(0.031)	(0.035)
Early career	0.044***	-0.058**
	(0.016)	(0.136)
Experience	0.024**	-0.038**
	(0.010)	(0.009)
Stocks covered	-0.001	0.004***
	(0.001)	(0.001)
Industries covered	0.009**	-0.003
	(0.004)	(0.04)
Times nominated	-0.007	-0.007
	(0.006)	(0.005)
Additional headline controls	Y	Y
Year fixed effects	Y	Y
Main industry covered fixed effects	Y	Y
Analyst-by-stock fixed effects	Y	Y
Employer fixed effects	Y	Y
First-stage F statistic	72.6	72.6
t -statistic on IV	8.5	8.5
Analyst-stock-quarter observations	594,129	594,129

Table B.8**Quality of Analysts' Reports: Variation Across Stocks Covered**

This table presents IV estimates of the relationship between AI intensity and the quality of analysts' reports at the analyst-stock-quarter level. In Columns (1)–(2), the dependent variables are analyst accuracy and bias as a percent of the mean consensus forecast among all analysts tracking stock i in quarter t . In Columns (3)–(4), the dependent variables are analyst accuracy and bias as a percent of the stock price at the close of the previous quarter. In Column (5), the dependent variable is the boldness of the analyst's forecast, a proxy for effort. The primary independent variable of interest is AI intensity which measures the quantity of social media posts that an AI-powered investment tool analyzes for stock i in quarter t . The IV for AI intensity is an indicator variable for whether the newspaper headlines in the USA Today for that stock were shorter than the median headline length in that quarter. Additional analyst controls include total number of stocks covered, total number of industries covered, general experience as an analyst, and reputation. Additional stock controls include the total key news events, total analyst coverage, firm size, average monthly returns, volatility, market-to-book ratio, volatility of ROE, Amihud's illiquidity ratio, profitability, an indicator for being in the S&P 500, and counts of key newspaper headline words. Below the coefficient estimates are robust standard errors clustered at the stock level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable =				
	As % of consensus Accuracy (1)	Bias (2)	As % of stock price Accuracy (3)	Bias (4)	Bold forecast (5)
AI intensity	-0.068* (0.038)	-0.014 (0.032)	-0.052*** (0.016)	0.011 (0.017)	0.064*** (0.018)
Additional analyst & stock controls	Y	Y	Y	Y	Y
Analyst-by-year fixed effects	Y	Y	Y	Y	Y
Main industry covered fixed effects	Y	Y	Y	Y	Y
Employer fixed effects	Y	Y	Y	Y	Y
First-stage F statistic	105.6	105.6	107.4	107.4	107.6
t -statistic on IV	10.3	10.3	10.4	10.4	10.4
Analyst-stock-quarter observations	600,002	600,002	631,453	631,453	603,995

Table B.9

Quality of Analysts' Reports: Quantile Regressions

This table presents quantile regressions estimates of the relationship between AI intensity and the accuracy of analysts' reports at the analyst-stock-quarter level. The focal independent variable is AI intensity which measures the quantity of social media posts analyzed to generate AI-powered stock picks for stock i in quarter t . Each column represents a different quantile of accuracy. Additional analyst controls include total number of stocks covered, total number of industries covered, general experience as an analyst, and reputation. Additional stock controls include the total key news events, total analyst coverage, firm size, average monthly returns, volatility, market-to-book ratio, volatility of ROE, Amihud's illiquidity ratio, profitability, an indicator for being in the S&P 500, and counts of key newspaper headline words. Below the coefficient estimates are robust standard errors clustered at the stock level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable =					
	Accuracy (as % of consensus)					
	15th	25th	Median	75th	85th	
	(1)	(2)	(3)	(4)	(5)	
AI intensity	-0.160***	-0.112	-0.055	-0.030	-0.020	
	(0.060)	(0.099)	(0.191)	(0.233)	(0.249)	
Additional analyst & stock controls	Y	Y	Y	Y	Y	
Analyst-by-stock fixed effects	Y	Y	Y	Y	Y	
Main industry covered fixed effects	Y	Y	Y	Y	Y	
Employer fixed effects	Y	Y	Y	Y	Y	
Year fixed effects	Y	Y	Y	Y	Y	Y
Analyst-stock-quarter observations	594,143	157,642	436,499	594,143	257,900	

Table B.10**Quality of Analysts' Reports: Heterogeneous Treatment**

This table presents IV estimates of the relationship between AI intensity and the quality of analysts' reports at the analyst-stock-quarter level. The dependent variable is analyst accuracy as a percent of the mean consensus forecast among all analysts tracking stock i in quarter t . In Column (1), the focal independent variables are AI intensity and the interaction between AI intensity and an analyst being in the top third of accuracy historically. In Columns (2) and (3), the focal independent variable is AI intensity but the samples are split into analysts historically accurate and not. The IV for AI intensity is an indicator variable for whether the newspaper headlines in the USA Today for that stock were shorter than the median headline length in that quarter. The IV for the interaction term is the same IV interacted with the indicator for being historically accurate. Additional analyst controls include total number of stocks covered, total number of industries covered, general experience as an analyst, and reputation. Additional stock controls include the total key news events, total analyst coverage, firm size, average monthly returns, volatility, market-to-book ratio, volatility of ROE, Amihud's illiquidity ratio, profitability, an indicator for being in the S&P 500, and counts of key newspaper headline words. Below the coefficient estimates are robust standard errors clustered at the stock level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Dependent variable = Accuracy (as % of cons.)	Dependent variable =		
	All	High Acc.	Non-High
	(1)	(2)	(3)
AI intensity	0.083 (0.078)	0.143* (0.075)	0.106 (0.082)
AI intensity \times Historically highly accurate	0.094** (0.047)		
Additional analyst & stock controls	Y	Y	Y
Analyst-by-stock fixed effects	Y	Y	Y
Main industry covered fixed effects	Y	Y	Y
Employer fixed effects	Y	Y	Y
Year fixed effects	Y	Y	Y
First-stage F statistic	30.8	68.6	54.0
Analyst-stock-quarter observations	594,143	257,895	336,231

Swiss Finance Institute

Swiss Finance Institute (SFI) is the national center for fundamental research, doctoral training, knowledge exchange, and continuing education in the fields of banking and finance. SFI's mission is to grow knowledge capital for the Swiss financial marketplace. Created in 2006 as a public-private partnership, SFI is a common initiative of the Swiss finance industry, leading Swiss universities, and the Swiss Confederation.