Artificial Intelligence and the Future of Work: Evidence from Analysts*

JILLIAN GRENNAN
Duke University
Fuqua School of Business

RONI MICHAELY
University of Geneva
Geneva Finance Research Institute

ABSTRACT

Artificial intelligence (AI) can enhance prediction, a common task for highly-skilled workers. We examine the implications of this technological change for incumbent workers in the context of security analysts. As evidence of substitution, we find the most talented analysts quit the profession while others shift their coverage toward low-AI stocks. Analysts' access to management gives them a soft information advantage, and they focus these meetings on low-AI stocks suggesting some complementarity. The quality of analysts' predictions also change: analysts covering high-AI stocks exhibit increased bias and forecast errors. Additional tests suggest the change in reporting quality stems from a less talented pool of analysts rather than strategically biased predictions.

JEL classification: D14, G11, G14, G23, O35

Keywords: Sell-side equity research analysts, Investment recommendations, Earnings estimates, Artificial intelligence, Market intelligence equity FinTechs, Financial technology (FinTech), Aggregators, Machine learning, Social media, Big Data, Labor Displacement, Productivity

*Authors: Grennan, Duke University (e-mail: jillian.grennan@duke.edu); Michaely, University of Geneva (e-mail: roni.michaely@unige.ch). We thank Jaeyeon Lee for excellent research assistance. Some of the data used in this study comes from TipRanks, a firm in which Roni Michaely has an equity interest and serves on the Board of Directors.

"With 80% of the data in the world created in the last two years, judgment matters more than ever. Technology is a complement to sound judgment and knowledge, not a substitute."

— Joyce Chang, Global Head of Research, J.P. Morgan, September 2017

I. Introduction

The implications of technological change for incumbent workers are a source of controversy (Levy and Murnane, 2004; Autor, 2015). Some scholars argue that the most recent wave of artificial intelligence (AI) is a forerunner to potential widespread joblessness, especially among high-skilled workers. Others argue that AI, like previous forms of technological change, will ultimately increase labor demand and enhance productivity. Still others argue the impact is ambiguous because AI companies can only automate a particular set of job tasks. To understand the controversy surrounding labor markets, it is important to place AI in context. Over the past decade, advances in computer speed, data gathering, and techniques for processing data previously considered too cumbersome and unwieldy to yield insights have substantially improved. These advances mean that AI algorithms, which are generally tasked with predicting some outcome, have become more accurate.

Prediction is essential to many occupations. For example, human resources professionals use resumes to predict which workers to recruit and hire but an AI algorithm can also be used to process data and produce such predictions. Similarly, security analysts gather information to predict changes in company fundamentals such as revenue or earnings that ultimately drive stock returns but AI can also be used to predict future earnings or stock returns. Given that such prediction is an important job task for highly-skilled workers, such as sell-side security analysts, some researchers argue that this wave of technological change is different than previous periods of technological change (Agrawal, Gans, and Goldfarb, 2019).

Theoretical literature on AI recognizes its importance as a prediction technology, and therefore, focuses on how the technology will impact incumbent workers in a task-based framework (Acemoglu and Restrepo, 2018; Brynjolfsson, Mitchell, and Rock, 2018). Nevertheless, the theoretical literature yields ambiguous answers when it comes to the potential for AI to serves as a substitute, complement, or some mix thereof for incumbent workers. One strand argues that AI may directly substitute for labor in prediction tasks or tasks that were previously not framed as a prediction tasks will be transformed into a prediction task. While this suggests AI will serve as a substitute to highly skilled labor, it is also the case that by removing the prediction task or increasing the reliability of the predictions, highly-skilled workers may be able to focus their attention on new or complementary tasks. In this sense, the other strand of the theoretical literature argues that AI will ultimately increase labor demand and productivity (Autor, Levy, and Murnane, 2003; Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2019). Finally, given that a new prediction itself discloses information, the response of workers to AI may be more nuanced. For example, workers may recognize how their input feeds back into the prediction algorithm, so they may strategically alter the inputs they are expected to provide (Bond, Edmans, and Goldstein, 2012; Goldstein and Yang, 2019).

In contrast to the theoretical side, empirical work on the effects of the AI technological wave on workers has been limited although insights can be drawn from some studies examining the replacement of labor by machines (Lin, 2011; Acemoglu and Restrepo, 2017; Graetz and Michaels, 2018). We aim to fill this gap by providing empirical evidence on the response of incumbent workers to AI in the financial services industry and we focus on sell-side equity analysts. This setting is an ideal one for studying this issue for a couple of reasons. First, the prediction problem faced by security analysts is well-defined and data is readily available documenting their predictions. Yet, their job is ripe for disruption by AI given that analysts' predictions are notoriously biased because of conflicts of interest (Michaely and Womack, 1999; Hong and Kacperczyk, 2010). Second, both of the more nuanced economic channels about task complementarity and feedback effects are likely to apply in the market for analysts' forecasts. Finally, there is meaningful variation across securities in terms of data available for AI to use in its prediction algorithm. From an econometric perspective, this creates an opportunity to exploit this variation for identification purposes.

In particular, we gain access to data gathered by TipRanks, a market intelligence financial technology firm or ("FinTech") operating in this space. These FinTechs are streamlining and synthesizing many data sources, including nontraditional ones, relevant for equity investment recommendations (Grennan and Michaely, 2019). As one director of equity research said about these FinTechs, "you can get the benefits of hundreds of stock pickers in a low-cost quantitative way," (Loder, 2019). We use their social media data as it is the most common sources of nontraditional data analyzed by these FinTechs and there is meaningful variation in the extent to which stocks are popular on social media. For example, Cookson and Niessner (2019) document a disproportionate amount of StockTwits are about Apple and Facebook. Specifically, TipRanks provided information on millions of social media and blog posts that they had linked to specific equities between 2010 and 2017. We proxy for AI and big data intensity in a given security using the number of posts that provide financial analysis on a particular stock. The details of this data show that it is very noisy. This suggests FinTechs ability to use AI algorithms to analyze these data rapidly to predict the quality of the information is an important feature.

To determine analysts' response to this type of AI and big data, we use an instrumental variable (IV) approach. This is necessary because there are likely to be factors that are unobservable econometrically that may be related to both AI and big data use and analysts responses. If this is the case, then tests using ordinary least squares (OLS) estimates would be difficult to interpret. Our IV approach uses the insight that social media influencers and bloggers follow what is popular, often only adding their own commentary. While popularity is not random, variation in popularity can come from things that are quasi-random. Clickbait or text designed to entice internet users to follow the link is one such example. Recent 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).

We build on this short title idea. Using data from RavenPack with all headlines from major newspapers, our IV is an indicator for an equity having newspaper headline lengths below the median headline length. Weak instrument tests show that this short-title instrument exhibits a statistically significant association with increased blogging at the 99th percentile, which is consistent with bloggers chasing popularity. This IV plausibly satisfies the exclusion restriction because newspaper headline length is assigned by an editor in a way that is independent of an article's content.[1] To further test if headline length is quasi-random, we regress headline length on firm characteristics for more than 7 million articles and find no evidence of selection. This also holds true when we use variable selection techniques to determine the words associated with headline length; the selected words appear to be random rather than systematically linked to firm characteristics. Finally, by controlling for total newspaper coverage, we ensure the IV captures incremental popularity differences rather than significant news regarding a firm driving AI intensity in a given stock.

Using the IV strategy, we begin by examining on the direct substitution hypothesis that AI leads to fewer analysts. We first examine analysts' employment and find evidence of substitution. In particular, we find that analysts are leaving the profession at a higher rate when they stocks they cover have high AI and big data intensity. Second, we examine analysts' decisions of which stocks to cover. We see that the analysts who remain in the job shift their coverage toward low AI and big data stocks. This shift in coverage is interesting

[1] In printed newspapers, column space is predetermined by the number of advertisements sold. Thus, in the Wall Street Journal print edition, the stories that make the front page must adhere to an assigned column width. The editors are responsible for generating a title that fits within the predetermined space limit.

because it could be either direct substitution or a more nuanced form of complementarity.

On one hand, the shift in coverage by analysts toward low AI and big data stocks could be evidence of AI substituting for high-skilled labor by reducing demand for analysts. On the other hand, the shift in coverage could represent a more nuanced response to AI and big data whereby it complements high-skilled labor. In this case, AI and big data improve and change the nature of sell side analysts' work. For example, this shift in coverage could occur because the enhanced predictive capabilities allow analysts to focus their efforts on the stocks where data is not readily available. For such stocks, meetings with management, their suppliers, and hosting investor conferences are ways in which analysts use their interpersonal skills, something that cannot be automated, to gather soft information. Beyond soft information, analysts also have soft skills, such devoting time to institutional clients to provide color for their recommendations. To disentangle the direct substitution hypothesis from the more nuanced complementarity hypothesis, we gather data from Bloomberg Corporate Events on meetings with management. We find that analysts are indeed focusing their time with management on low AI and big data stocks, suggesting some degree of complementarity.

Next, we focus on the quality of the output of the traditional task of analysts, their earnings forecast. We examine the accuracy and the optimism bias in their earnings forecasts. We find that overall reporting quality declines. We find higher aggregate absolute forecast errors and more optimistic bias for stocks where big data and AI intensity is likely to be higher. A one standard deviation increase in AI and big data intensity is associated with a 0.14 standard deviation increase in aggregate optimistism bias and a 0.23 standard deviation decrease in forecast accuracy. To put this number in context, we compare our point estimate to that of other standard covariates in the literature (i.e., analyst coverage, firm size, recent stock returns) and find that AI and big data intensity rank in the middle relative to other variables used to explain analysts' reporting quality.

Next, we assess the economic mechanism through which AI and big data intensity may influence the reporting quality. In particular, we assess whether the bias is driven by the effort analysts devote, their decision to switch careers, and/or the extent to which they bias their recommendations for strategic reasons. For example, if lower reporting quality stems from a change in the talent pool for analysts, such a mechanism could offset potential benefits from AI and big data. Our main finding is that the best analysts are quitting and leaving the profession, which suggests a permanent reduction in the quality of analysts' reports rather than a strategic response.

The key takeaways of our paper is that the relation between AI and the future of high-skilled work is nuanced. We find evidence of direct substitution via job loss. We see that analysts are more likely to leave the profession when they cover stocks where big data and AI intensity is likely to be higher. But we also find evidence of a more nuanced complementarity. AI and big data are freeing up analysts' time to use their soft skills – ones that cannot be automated – to engage with management and clients. The additional time for soft skills only improves reporting quality for a small number of stocks, so on average, reporting quality decreases because the soft skill gains cannot offset the loss of talent. Overall, the evidence from sell-side analysts suggests the AI and big data wave of technological change is different than previous periods.

On a broad level, this paper contributes to the economic literature on labor market responses to technological change. It is related to the empirical literature on the effects of technology on wages (Katz and Murphy, 1992), employment (Autor, Levy, and Murnane, 2003; Autor and Dorn, 2013; Autor, Dorn, and Hanson, 2013; Michaels, Natraj, and Van Reenen, 2014), and the theory of the firm given how technology by change the boundary of human capital (Bartel, Ichniowski, and Shaw, 2007; Buchak, 2019). By focusing on AI technology, we relate to work examining its economic consequences (Brynjolfsson, Mitchell, and Rock, 2018; Agrawal, Gans, and Goldfarb, 2018; Erel, Stern, Tan, and Weisbach, 2018). This paper also contributes to our understanding of the market for security analysis. Our evidence suggests that, on average, the data being processed by market intelligence equity FinTechs is crowding-out private information production by decreasing the reporting quality of analysts. This finding parallels the information disclosure literature (Bond, Edmans, and Goldstein, 2012; Goldstein and Yang, 2019) in that both generate crowd-out effects. This finding runs counter to research on competition and bias (Hong and Kacperczyk, 2010). That AI-based FinTechs do not produce the same competition-related incentives for analysts that the entry of more analysts do is novel and overall, these findings help to explain the economic forces shaping analysts' recommendations (Hong and Kubik, 2003; Barber, Lehavy, and Trueman, 2007; Fang and Yasuda, 2009; Merkley, Michaely, and Pacelli, 2017).

II. Data

A. Sell-side Equity Analysts

We use comprehensive analyst earnings forecast and recommendation data from IBES to derive our main measures of analyst reporting quality. We supplement these data with data from Zacks in some analysis as IBES stopped providing broker or analyst identity in 2007. We use share price data from the Center for Research in Security Prices (CRSP) and accounting data come from the CRSP-Compustat merged database. Supplementary sources of data include equity ownership data from Thomson-Reuters, and mergers and acquisition (M&A) and securities issuance data from SDC. Our main dependent variables of interest are the bias and accuracy in the consensus analyst forecast.

To calculate quality measures of analysts' reports, we compute signed forecast errors for each analyst-firm forecast as the difference between the consensus earnings per share (EPS) forecast minus actual EPS, scaled by the absolute value of the consensus EPS forecast, such that positive forecast error indicate greater optimism bias. Our consensus number is taken from Compustat due to data issues with IBES (Ljnuqvist, Malloy, and Marston, 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. We compute absolute forecast error as the absolute value of signed forecast error. Intuitively, aggregate absolute forecast error is a proxy for accuracy. Again, similar to prior studies that use aggregate measures, we equal-weight forecast errors. This procedure essentially weights all analyst forecasts equally so as not to obscure any individual forecast.

Analysts' revisions to their earnings estimates and recommendations are associated with many other factors. As such, our main control variables include those that are standard in the literature. For our primary specification the main controls are newspaper coverage, analyst coverage, firm size, daily return volatility, mean monthly returns, market-to-book ratio, volatility of return on equity (ROE), profitability, and membership in the S&P 500. All of these controls, with the exception of newspaper coverage, match those used by Hong and Kacperczyk (2010) to study the effect of analyst supply on bias. We add newspaper coverage to our main analysis for two reasons. First, it helps ensure our instrument headline length is not simply serving as a noisy proxy for newspaper coverage. Second, research suggests that greater newspaper coverage is associated with both greater analyst coverage and lower analyst bias.

When using data at the analyst level as opposed to the equity level, we include the following analyst-level control variables: general experience, firm experience, firms covered, industries covered, forecast frequency, forecast horizon, days since last forecast, employer affiliated with firm, brokerage size, and independent broker status. These controls reflect insights from a broad literature showing the importance of analysts' career paths, potential conflicts of interest, and the mitigating role of institutions, herding, reputation, communication with insiders, and industry expertise.

B. Bloomberg Data on Analyst-Management Meetings

Sell-side equity analysts often organize conferences and meetings with management. We obtain data on broker-hosted conferences from the Bloomberg Corporate Events Database. The database includes information on conference name, date, hosting brokerage firm, presenting company names and ticker symbols. We then match the firms participating in the conferences in a given quarter to specific stocks and the analysts that work at the brokerage firms hosting the events.

C. AI-based Stock Predictions

Market intelligence equity FinTechs streamline and synthesize many data sources, including nontraditional ones, relevant for equity investment recommendations (Grennan and Michaely, 2019). These FinTechs analyze nontraditional data sources previously considered too cumbersome or inconsistent to inform investment decisions. To approximate the competition analysts face from these market intelligence equity FinTechs, we gain access to data gathered by TipRanks, a FinTech operating in this space. We focus on their social media data as it is the most common sources of nontraditional data analyzed by market intelligence equity FinTechs and there is meaningful variation across stocks in terms of social media coverage. For example, It is also the case that hedge funds and other investors frequently pay high fees to gain access to buy and sell signals that these FinTech derive from social media data. For example, Cookson and Niessner (2019) document a disproportionate amount of StockTwits are about Apple and Facebook. Specifically, TipRanks provided information on millions of social media and blog posts that they had linked to specific equities between 2010 and 2017. We use the number of posts that provide financial analysis on a particular stock as our proxy for the degree of AI-based competition or AI intensity for a given stock in a given quarter. While this social media data is not representative of all nontraditional data sources that encompass big data which AI-based FinTechs may be using to predict stock price movements, it does provide insight as a proxy for the competition security analysts face from such automation.

D. Summary Statistics

Table 1 provides descriptive statistics about the key dependent and independent variables used in this study. We winsorize all variables at the 1st and 99th percentiles to minimize the influence of outliers. The formulas for the variables derived from these databases are included in Appendix A. Due to the merging of various datasets, our main sample period runs from 2010Q1 to 2016Q3, resulting in 81,597 observations at the equity-quarter level.

III. Hypotheses Development

Given the claims from prior research yield ambiguous answers when it comes to the potential for AI to serves as a substitute, complement, or some mix thereof for incumbent workers, we examine each hypothesis separately. The specific questions we ask to test each hypothesis are briefly outlined below.

- **Substitution hypothesis**: Are analysts changing the stocks they cover or switching careers in response to AI?
- **Complements hypothesis:** Do the enhanced predictive capabilities of AI free up time for analysts to focus on their competitive advantage (i.e., the gathering of soft information)?
- **More nuanced economical channels for incumbent workers response:** Do analysts change the quality of the work product they provide (i.e., generating an earnings forecast)? If so, what economic channel is driving this change? Is it the effort analysts

devote, their decision to switch careers, and/or the extent to which they bias their recommendations for strategic reasons?

IV. Empirical Strategy

To test the hypothesis that analysts respond to the competition from AI, we use data from a market intelligence equity FinTechs. The market intelligence FinTech is recognized as a leader in the AI space having won awards for best FinTech and having its technology incorporated into popular trading platforms such as E-Trade Financial. Using this data, we first study changes in stock coverage by analysts and their decisions to exit the profession. Next, we study specific activities of analysts such as their decisions to engage in meetings with management. Then, we examine the quality of the work done by analysts. Finally, we examine the economic channels that may be generating changes in the quality of the work done by analysts. In each case, we use an IV strategy. Such a strategy is advantageous because tests that rely on ordinary least squares (OLS) regressions are difficult to interpret due to endogeneity concerns. For example, if analyst bias is correlated with consumer popularity, which is unobservable econometrically, then the point estimate in the OLS regression is likely biased by this unobservable factor. In practice, there is likely to be more than one unobservable factor correlated with AI and big data intensity, and each factor is likely to push the bias to be either positive or negative. As such, it is difficult to determine the net effect of these unobservable factors (Roberts and Whited, 2013).

To provide a credible point estimate and mitigate the influence of factors endogenous to the data-generating process for analysts' reporting quality, we use an IV approach. Our IV is an indicator for whether the equity has below median headline length in a given quarter. The intuition for why this instrument is relevant is that it generates variation in the frequency of nontraditional sources of information like financial blogs or social media writing about that equity in a given quarter. Nontraditional information producers such as financial bloggers are more likely to click on a newspaper article about an equity when it has a short headline. Given that they clicked on the article, they are also likely to be inspired to write their own commentary about such a stock.

More specifically, the relevance condition is that financial bloggers focus on what is popular and thus are more likely to write a blog about an equity after reading an article about that equity. The exclusion restriction for IV identification requires that the short-headline newspaper coverage only alter analysts' aggregate accuracy or bias via the effect of additional blogging about an equity increasing the concentration of FinTechs in that equity. The main argument supporting the plausibility of the instrument is that financial news headlines are quasi-random because they are selected at the discretion of the editor.

To construct a quarterly headline length indicator, we obtain newspaper headline data from RavenPack. We include in our analysis only the following newspapers, which have the highest national readership in the United States: USA Today, the Wall Street Journal, the New York Times, the Los Angeles Times, the Chicago Tribune, the Washington Post, the Financial Times, and the DowJones Newswire. As reported in Appendix Table B1, both the mean and the median headline length is 57 characters, while the 25th percentile and the 75th percentile are 48 and 63 words, respectively.

Appendix Table B2 provides several example headlines for the firm Apple when the length is at the 25th percentile and the 75th percentile. The example headlines show that there is no discernable difference in content conveyed by the length of the headline. To provide additional evidence that headline length is quasi-random, Appendix Table B3 provides the results of a regression in which headline length is the dependent variable and firm characteristics such as market-to-book ratio, profitability, ROE, momentum, and firm size are the explanatory variables. No variable is statistically significant at the 95th percentile. Moreover, the R-squared from a regression with more than 7 million observations is only 0.10%.

In contrast, studies that have looked at the sentiment of financial news in relation to these exact same variables show highly significant correlations with these characteristics (Niessner and So (2017)).

As a final argument for the plausibility of the IV, we employ the model selection technique of LASSO (Efron, Hastie, Johnstone, and Tibshirani (2004)) to determine whether the words associated with title 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 "mgmt" are associated with shorter headlines. In contrast, if we thought the instrument was erroneously capturing sentiment or conveying content, we would expect to see words like "beat" as in "beat earnings estimates" or "fear" as in "fear trade conflicts," but we do not see any of these words.

Finally, because we control for total newspaper coverage in the first stage of the IV regressions, the instrumented value for AI intensity is being identified from the incremental effect of having short headlines over and above having any headline. Hence, even if analysts also read newspaper articles about the equities they cover, which is likely given that being current on the latest developments is exactly what analysts are supposed to do, this is not the incremental variation we are using. Rather, we are relying on shorter headlines increasing clicks, an important metric in the era of internet search engine optimization.

Our exact IV specification is as follows:

$$Outcome_{it} = \alpha + \beta AI_{it} + \theta X_{it} + f_i + \delta_t + \epsilon_{it} \tag{1}$$

where $Outcome_{it}$ captures the analysts' decisions about which stocks to cover and or their reporting quality in terms of optimism bias and accuracy in quarter t for equity i. AI_{it} is our proxy for AI intensity that is based on the quantity of alternative data sources such as financial blog posts that are processed by market intelligence FinTechs and that discuss equity i in quarter t. X_{it} is a vector of observables (newspaper coverage, analyst coverage, firm size, daily return volatility, mean monthly return, log market-to-book ratio, volatility of ROE, profitability, and an indicator for if the stock is a member of the S&P 500). f_i is an firm fixed effect; δ_t is a quarter fixed effect; and ϵ_{it} is the unobservable error component.

When performing the analyses, we rely on cross-sectional variation because the point estimate based off of cross-sectional variation captures the full range of AI intensity, from no alternative information processed to very high quantities of nontraditional information processed. In contrast, within-firm estimates likely only capture small variations ranging from no to low coverage or high to very high coverage. This is due to the fact that once data is available for AI-based FinTechs to analyze, the level of coverage a firm receives is fairly persistent.

To disentangle the economic channels driving analysts' actions, we also examine the market's response to changes in analysts' recommendations by analyzing excess returns and excess volume associated with a recommendation change. 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, and we measure it relative to the recommendation announcement at day 0. 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:

$$CAR_{ijt} = \alpha_{ijt} + \beta AI_{ijt} + a_j + \delta_t + e_{ijt} \tag{2}$$

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

V. Results

In this section, we explore whether analysts respond to the big data and AI coverage that is available for certain equities. We begin by examining a substitution hypothesis. Table 2 shows that when AI and big data intensity for a stock increases, the analysts that cover that stock are more likely to quit and leave the profession. When we further isolate the variation in AI and big data intensity, we see the result is entirely driven by stocks with increasing AI and big data coverage. These findings are statistically significant at the 99th percentile. Weak instrument tests suggest that our instrument is relevant, and thus, producing a consistent estimate. In Table 3, we examine an analogous hypothesis for the analysts that stay. Specifically, we test whether the analysts who stay in the profession change the stocks that they make forecasts about. Again, we see that when AI and big data intensity for a stock increases, the analysts both terminate coverage of the stock and fail to initiate new coverage in stocks with high-AI intensity. These findings are statistically significant at the 99th percentile.

This finding that AI and big data capabilities are a direct substitute for the primary job tasks of security analysts is economically meaningful. It suggests that the scope of the analyst industry is changing in response. Such findings are consistent with career concerns models such as those posited by Bar-Isaac (2005) in which a non-monotonic relation exists between competition and reputation. In the context of analysts, the idea is that as competition increases from AI, at some point an analysts' reputations will no longer be a concern as they know their industry is changing; they care less about their reputation in that industry as they plan to leave it soon. Thus, FinTechs could encourage some more experienced analysts to leave their firms due to their reduced position of prominence and prestige. Further, as the marginal value of information production by analysts decreases (due to market intelligence equity FinTechs entrance), analysts compensation is likely to decrease, resulting in lower incentives to stay. AI, therefore, could change the overall composition of who chooses to be sell-side analysts (Merkley et al. (2017)). If this were to occur, the pool of analysts would be younger and less qualified compared to previous generations. In other cases, an analysts' professional relationships could determine whether he or she chooses to stay. For example, an unaffiliated analyst may be more likely to leave the industry than affiliated analysts. We explore these economic channels in more detail.

One potential way the scope of the analyst industry may change is by allowing analysts to focus their efforts on stocks for which less information is readily available. To gain insights into such stocks, analysts often need to arrange meetings with management to try and garner soft information from them. Table 4 shows that analysts do focus their efforts in gathering information about stocks with lower AI-intensity. This is consistent with the task-based theory of how technology will change work. It suggests that AI will serve as a replacement for easier prediction problems, but this then will free up time for employees to work on harder tasks such as predicting the stock price movements of stocks that require gathering soft information first. This suggests at least in the context of analysts there is some complementarity from AI.

Our next set of analysis examines the quality of the product produced by analysts. As is reported in Table 5, we find that when AI intensity increases, analysts' consensus optimism bias increases but their consensus accuracy declines. In comparison to the 0.05 estimated increase in optimism bias from the OLS estimation, the point estimate is 0.14, as reported in Column (1). This doubling of the point estimate is consistent with many potential unobservable variables biasing the OLS estimate and, in aggregate, pushing the uninstrumented point estimate toward zero. Column (3) shows that a one standard deviation increase in FinTech coverage is associated with a 0.23 standard deviation decrease in mean aggregate accuracy.

The statistical evidence for the deterioration in reporting quality is significant at the 99th percentile. The F-statistic from the first stage of the instrumental variable regression is 195.9, which exceeds the requisite 10 to ensure minimal bias of the point estimate. The t-statistic on the instrument in the first-stage is 14.0, which suggests the instrument is not weak. The IV specification includes controls for newspaper coverage, analyst coverage, firm size, daily return volatility, mean monthly returns, market-to-book ratio, volatility of ROE, profitability, and membership in the S&P 500. These controls help to account for other firm-level dynamics that could lead to a deterioration in analyst reporting quality.

The details of the point estimates on the additional controls convey useful information. First, AI intensity is not changing any traditional relationships observed in the data. For example, increased analyst coverage and greater average monthly stock returns are both still associated with lower bias Hong and Kacperczyk (2010). Second, our results confirm recent research suggesting that increased newspaper coverage is associated with reduced bias (Bradshaw, Wang, and Zhou, 2017). Third, the results help put FinTech coverage in context by showing where AI intensity ranks relative to other firm-specific 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, return volatility, and profitability but larger than that of newspaper coverage, market-to-book ratio, and monthly returns.

Next, we assess the economic channel inducing analysts to change their reporting quality. If the lower reporting quality stems from a change in the talent pool for analysts, this would be difficult to reverse. In contrast, if the lower reporting quality stems from low effort, contractual mechanisms could more easily be introduced to reverse the effect. Table 6 presents evidence testing the hypothesis that AI intensity is having an anti-competitive effect on analysts' reporting quality via career concerns. Panel A presents evidence supporting a career concerns channel. As reported in Column (1) the analysts who leave the profession are among the top 25% most accurate analysts, respectively. The remaining regressions analyze the same dependent variables but categorize the independent variable, FinTech coverage, into growing coverage and no change in coverage. For each outcome, growing coverage versus no change in coverage produce asymmetric responses. For stocks with growing FinTech coverage, 1.6 more analysts quit and leave the profession, while 3.2 fewer analysts leave the profession when there is no change in FinTech coverage. All of the results are statistically significant at the 95th or 99th percentile, but the instrument is weaker when the independent variable is no change in FinTech coverage.

Table 6, Panel B tests another channel through which AI intensity could affect analyst coverage: analysts generating bolder forecasts to garner publicity or attention. 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. We follow the prior literature and classify forecasts as bold if they are either above or below both the analyst's own prior forecast and the consensus forecast immediately prior to the analyst's forecast. In Columns (1) – (6), we examine both the percentage of bold forecasts made and the distance the forecast is from the consensus forecast. We find no evidence to suggest that analysts are strategically responding to AI intensity by generating more audacious forecasts. In fact, we find the opposite. As in Panel A, when we examine growing coverage versus no change in coverage, we find an asymmetric response. Thus, this test is consistent with analysts exerting less effort. Addressing an effort problem can likely be fixed faster than having to recruit a more talented pool of analysts.

Finally, Table 7 presents the results of tests of market responsiveness to analysts' recommendation revisions as a function of AI intensity. 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 when AI coverage is high. 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 equity is high. Both results indicate that when FinTechs with AI capabilities are predicting stock prices, the information produced in analysts' reports is, increasingly, less relevant and possibly already impounded in prices.

Panel B of Table 7 uses an alternative approach to identify the market responsiveness to analysts when AI intensity is high. In Panel A, the associations come from variation in the cross-section of stocks covered by analysts. Yet a comparison of the market's reaction to analyst recommendations for the same firm before and after FinTech entry is also prudent. Because our data set does not contain the exact date upon which different FinTechs started covering different equities, we instead compare mean changes in market responsiveness over multiple years. Specifically, we run a placebo-like test that compares the five years before our sample starts to the years of data in our sample. We add the analyst recommendation changes in the five years prior (2005–2009) and set the FinTech coverage to 0. We then focus on the within-analyst or within-analyst-equity variation. Thus, our coefficient of interest is similar to the coefficient in a triple difference-in-differences estimator.

As reported in Columns (1) and (2) of Panel B, the estimates indicate that the market is 28 and 25 basis points less responsive to analysts' recommendation revisions for the same firm over time when that firm has greater AI intensity. These results are significant at the 99th and 90th percentile, respectively. Columns (3) and (4) of Panel B run the same test using the earlier data and reveal a negative point estimate that is significant at the 90th percentile. Overall, both the cross-sectional estimates in Panel A and the more restrictive, within-analyst-firm estimates in Panel B are consistent with the view that the marginal impact of the information analysts convey to the market is being reduced as market intelligence equity FinTechs help to incorporate nontraditional sources of data and predict stock price movements.

VI. 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 analyst earnings forecasts.

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 the most talented analysts quit the profession while others shift their coverage toward low-AI stocks. Analysts' access to management gives them a soft information advantage, and they focus these meetings on low-AI stocks suggesting some complementarity. The quality of analysts' predictions also change: analysts covering high-AI stocks exhibit increased bias and forecast errors.

Our analyses of what drives analysts' responses suggest that the drop-off in their reporting quality will not be easily remedied. The best analysts are quitting and leaving the profession, which suggests a permanent reduction in the quality of analysts' reports rather than a strategic response. Thus, the information presented to investors by market intelligence equity FinTechs is likely to play an increasingly important economic role in investors' decisions. This change, in turn, introduces new challenges as regulators learn how to best provide oversight of big data providers and the algorithms they use. For example, the financial advice provided by market intelligence equity FinTechs is not regulated, yet we know there is meaningful misconduct by humans in this area. Given that the algorithms are programmed by humans, they could have built-in bias.

REFERENCES

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., 2017. Robots and jobs: Evidence from us labor markets. Working Paper 23285, National Bureau of Economic Research.

Acemoglu, D., Restrepo, P., 2018. 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.

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.

Autor, D. H., 2015. Why are there still so many jobs? the history and future of workplace automation. Journal of Economic Perspectives 29, 3–30.

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., 2003. The skill content of recent technological change: An empirical exploration. Quarterly Journal of Economics 118, 1279–1333.

Bar-Isaac, H., 2005. Imperfect competition and reputational commitment. Economics Letters 89, 167–173.

Barber, B. M., Lehavy, R., Trueman, B., 2007. Comparing the stock recommendation performance of investment banks and independent research firms. Journal of Financial Economics 85, 490–517.

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.

Bond, P., Edmans, A., Goldstein, I., 2012. The real effects of financial markets. Annual Review of Financial Economics 4, 339–360.

Bradshaw, M. T., Wang, X., Zhou, D., 2017. Soft information in the financial press and analysts recommendation revisions., unpublished manuscript.

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.

Buchak, G., 2019. Financing the gig economy, unpublished manuscript.

Clement, M., Tse, S., 2005. Financial analyst characteristics and herding behavior in forecasting. The Journal of Finance 60, 307–341.

Cookson, A. J., Niessner, M., 2019. Why don't we agree? Evidence from a social network of investors. Journal of Finance forthcoming.

Efron, B., Hastie, T., Johnstone, I., Tibshirani, R., 2004. Least angle regression. The Annals of Statistics 32, 407–499.

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.

Goldstein, I., Yang, L., 2019. Good disclosure, bad disclosure. Journal of Financial Economics 131, 118–138.

Graetz, G., Michaels, G., 2018. Robots at work. The Review of Economics and Statistics 100, 753–768.

Grennan, J., Michaely, R., 2019. Fintechs and the market for financial analysis, unpublished manuscript.

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.

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.

Konnikova, M., 2014. How headlines change the way we think. New Yorker December 17.

Levy, F., Murnane, R. J., 2004. The New Division of Labor: How Computers Are Creating the Next Job Market. Princeton University Press.

Lin, J., 2011. Technological adaptation, cities, and new work. Review of Economics and Statistics 93, 554–574.

Ljnuqvist, A., Malloy, C., Marston, F., 2009. Rewriting history. The Journal of Finance 64, 1935–1960.

Loder, A., 2019. Goldman rolls out new etfs focused on articial intelligence. Wall Street Journal March 7.

Merkley, K., Michaely, R., Pacelli, J., 2017. Does the scope of the sell-side analyst industry matter? uppercaseAn examination of bias, accuracy, and information content of analyst reports. Journal of Finance 72, 1285–1334.

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., Womack, K. L., 1999. Conflict of interest and the credibility of underwriter analyst recommendations. Review of Financial Studies 12, 653–686.

Niessner, M., So, E. C., 2017. Bad news bearers: The negative tilt of financial press, unpublished working paper.

Roberts, M. R., Whited, T. M., 2013. Endogeneity in Empirical Corporate Finance, Elsevier, vol. 2 of *Handbook of the Economics of Finance*, chap. 0, pp. 493–572.

Table 1. Summary Statistics

This table presents summary statistics for the main dependent and independent variables. After combining the datasets, the main sample period is limited to 2010Q1 to 2016Q3. For a detailed description of each variable, see the definitions in Appendix A.

	Mean	Median	Std. Dev.	Obs.
	(1)	(2)	(3)	(4)
Dependent variables				
Analyst Quits and Leaves the Profession	0.16	0.00	0.42	81,597
Analyst Quits and Is Top 10% of Accuracy	0.01	0.00	0.11	81,597
Analyst Quits and Is Top 25% of Accuracy	0.02	0.00	0.14	81,597
Analysts Initiates Coverage on New Stock	0.04	0.00	0.26	81,597
Analysts Quits Covering a Stock	0.05	0.00	0.32	81,597
Analyst Meetings with Management	0.37	0.00	0.67	81,597
Analysts' Mean Bias (As % of the Abs. Value of Cons. EPS)	38.6%	6.7%	80.0%	81,597
Analysts' Median Bias	37.5%	4.4%	84.0%	81,597
Analysts' Mean Accuracy	60.9%	24.6%	77.1%	81,597
Analysts' Median Accuracy	57.2%	19.6%	80.5%	81,597
Bold Revision	0.65	0.66	0.19	81,597
Distance from Consensus	0.37	0.35	0.16	81,597
General Experience (Avg. Years of Experience)	4.26	4.18	1.69	81,597
Equity price informativeness ratio for analysts	0.07	0.05	0.07	81,597
Independent variables				
Artificial Intelligence (AI) Intensity	10.1	4.0	17.2	81,597
High-quality AI Predictions	2.7	1.0	6.0	81,597
Newspaper Coverage	81.1	10.0	666.7	81,597
Analyst Coverage	7.0	5.3	5.7	81,597
Firm size	13.9	13.9	1.7	81,597
Daily return volatility	37.9%	32.9%	20.6%	81,597
Mean monthly return	1.1%	1.2%	6.6%	81,597
Log market-to-book ratio	0.8	0.8	0.4	81,597
Volatility of ROE	98.8%	0.2%	685.1%	81,597
Profitability	1.8%	2.3%	4.5%	81,597
Member of S&P 500	15.0%	0.0%	35.7%	81,597
ROE	1.0%	2.2%	13.4%	81,597
Momentum	3.5%	0.6%	15.1%	81,597
Institutional ownership	59.8%	67.1%	30.6%	81,597
Hedge fund ownership	5.0%	0.6%	8.1%	81,597

Table 2. Career Choices of Analysts

This table presents IV estimates of the relationship between AI intensity and analysts' decisions to switch careers at the equity-quarter level. In Columns (1) – (3), the dependent variable is an indicator variable for if an analyst tracking stock i in quarter t quits the profession. The primary independent variable of interest, AI intensity, measures the quantity of financial blog posts analyzed by market intelligence equity FinTechs in quarter t that discuss equity i. Other independent variables of interest are an indicator for if AI intensity grows and an indicator for no change in AI intensity. The instrument for AI intensity is headline length, which indicates whether newspaper headlines about equity i in quarter t were shorter than the median headline length. Below the coefficient estimates are robust standard errors clustered at the equity level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively. For a detailed description of each variable, see the definitions in Appendix A.

	Dep. Var. = Analyst Quits		
	(1)	(2)	(3)
Artifical Intelligence (AI) Intensity	0.114***		
	(0.023)		
Growing AI Intensity		1.626***	
		(0.466)	
No Change in AI Intensity			-3.216***
			(1.062)
Additional Controls	Y	Y	Y
Time Fixed Effects	Y	Y	Y
First Stage F-Stat	195.9	35.5	15.7
T-Stat on Instrument	14.0	2.7	2.4
Adjusted R-squared	11%	10%	10%
Observations	81,597	63,544	63,544

Table 3. Stock Coverage by Analysts

This table presents IV estimates of the relationship between AI intensity and analysts' decisions to change the stocks they cover at the equity-quarter level. In Panel A, the dependent variable is an indicator variable for if an analyst initiates coverage of stock i in quarter t. In Panel A, the dependent variable is an indicator variable for if an analyst quits covering stock i in quarter t. The primary independent variable of interest, AI intensity, measures the quantity of financial blog posts analyzed by market intelligence equity FinTechs in quarter t that discuss equity i. Other independent variables of interest are an indicator for if AI intensity grows and an indicator for no change in AI intensity. The instrument for AI intensity is headline length, which indicates whether newspaper headlines about equity i in quarter t were shorter than the median headline length. Below the coefficient estimates are robust standard errors clustered at the equity level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively. For a detailed description of each variable, see the definitions in Appendix A.

	Dep. Var. = Analyst Initiates Coverage		
Panel A.	(1)	(2)	(3)
Artifical Intelligence (AI) Intensity	-0.039**		
	(0.019)		
Growing AI Intensity		-0.737***	
		(0.255)	
No Change in AI Intensity			-0.021
			(0.511)
Additional Controls	Y	Y	Y
Time Fixed Effects	Y	Y	Y
First Stage F-Stat	195.9	35.5	15.7
T-Stat on Instrument	14.0	2.7	2.4
Adjusted R-squared	11%	10%	10%
Observations	81,597	63,544	63,544

	Dep. Var. = Analyst Ends Coverage		
Panel B.	(1)	(2)	(3)
Artifical Intelligence (AI) Intensity	0.045**		
	(0.022)		
Growing AI Intensity		0.683**	
		(0.294)	
No Change in AI Intensity			0.034
			(0.428)
Additional Controls	Y	Y	Y
Time Fixed Effects	Y	Y	Y
First Stage F-Stat	195.9	35.5	15.7
T-Stat on Instrument	14.0	2.7	2.4
Adjusted R-squared	19%	22%	22%
Observations	81,597	63,544	63,544

Table 4. Analysts' Meetings with Management

This table presents IV estimates of the relationship between AI intensity and analysts' decisions to focus their efforts on their competitive advantage at gathering soft information by arranging meetings with management at the equity-quarter level. The dependent variable is the total number of meetings hosted by an analysts' brokerage firm where representatives for stock i in quarter t attended. The primary independent variable of interest, AI intensity, measures the quantity of financial blog posts analyzed by market intelligence equity FinTechs in quarter t that discuss equity i. Other independent variables of interest are an indicator for if AI intensity grows and an indicator for no change in AI intensity. The instrument for AI intensity is headline length, which indicates whether newspaper headlines about equity i in quarter t were shorter than the median headline length. Below the coefficient estimates are robust standard errors clustered at the equity level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively. For a detailed description of each variable, see the definitions in Appendix A.

	Dep. Var. = Meetings with Management		
	(1)	(2)	(3)
Artifical Intelligence (AI) Intensity	-0.174***		
	(0.027)		
Growing AI Intensity		-0.544***	
		(0.181)	
No Change in AI Intensity			0.045
			(0.211)
Additional Controls	Y	Y	Y
Time Fixed Effects	Y	Y	Y
First Stage F-Stat	195.9	35.5	15.7
T-Stat on Instrument	14.0	2.7	2.4
Adjusted R-squared	26%	24%	24%
Observations	81,597	63,544	63,544

Table 5. The Quality of Analysts' Reports

This table presents IV estimates of the relationship between AI intensity and the quality of analysts' reports at the equity-quarter level. In Columns (1) – (2), the dependent variable is analyst bias, defined as a consensus forecast bias of all analysts tracking stock i in quarter t. Forecast bias is the difference between the forecast of analyst j in quarter t and the actual EPS, expressed as a percentage of the consensus EPS. The consensus is obtained either as a mean as in Column (1) or a median as in Column (2). In Columns (3) – (4), the dependent variable is analyst accuracy, defined as a consensus absolute forecast error of all analysts tracking stock i in quarter t. The primary independent variable of interest, AI intensity, measures the quantity of financial blog posts analyzed by market intelligence equity FinTechs in quarter t that discuss equity i. Other independent variables of interest are an indicator for if AI intensity grows and an indicator for no change in AI intensity. The instrument for AI intensity is headline length, which indicates whether newspaper headlines about equity i in quarter t were shorter than the median headline length. Below the coefficient estimates are robust standard errors clustered at the equity level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively. For a detailed description of each variable, see the definitions in Appendix A.

	Bias (As % of EPS)		Accuracy (As % of EPS)	
	Mean	Median	Mean	Median
	(1)	(2)	(3)	(4)
Artifical Intelligence (AI) Intensity	0.14***	0.14***	-0.23***	-0.22***
	(0.04)	(0.04)	(0.04)	(0.04)
Newspaper Coverage	-0.02***	-0.02**	0.03***	0.03***
	(0.01)	(0.01)	(0.01)	(0.01)
Analyst Coverage	-0.01	-0.01	0.08***	0.07***
	(0.01)	(0.01)	(0.01)	(0.01)
Firm Size	-0.22***	-0.22***	0.24***	0.24***
	(0.02)	(0.02)	(0.02)	(0.02)
Daily Return Volatility	0.31***	0.30***	-0.33***	-0.32***
	(0.02)	(0.02)	(0.02)	(0.02)
Mean Monthly Return	-0.09***	-0.08***	0.06***	0.06***
	(0.00)	(0.00)	(0.00)	(0.00)
Log Market-to-Book	0.11***	0.11***	-0.07***	-0.07***
	(0.01)	(0.01)	(0.01)	(0.01)
Volatility of ROE	0.00	0.00	-0.00	-0.00
	(0.01)	(0.01)	(0.01)	(0.01)
Profitability	-0.39***	-0.37***	0.38***	0.36***
	(0.01)	(0.01)	(0.01)	(0.01)
Member of S&P 500	0.00	0.01	0.02**	0.01
	(0.01)	(0.01)	(0.01)	(0.01)
Time Fixed Effects	Y	Y	Y	Y
First Stage F-Stat	195.9	195.9	195.9	195.9
T-Stat on Instrument	14.0	14.0	14.0	14.0
Adjusted R-squared	38.7%	36.1%	40.6%	38.3%
Observations	81,597	81,597	81,597	81,597

Table 6. Economic Channel Behind Changes in Quality

This table presents IV estimates of the relationship between AI intensity and economic channels driving analysts to change their reporting quality at the equity-quarter level. Panel A examines a talent pool channel. Columns (1) – (3) examine whether more accurate analysts quit and leave the profession, and Columns (4) – (6) examine analyst experience. Panel B examines strategic responses related to career concerns and reputation. Columns (1) – (6) examine bold forecasts by analysts. The primary independent variable of interest, AI intensity, measures the quantity of financial blog posts analyzed by market intelligence equity FinTechs in quarter t that discuss equity i. Other independent variables of interest are an indicator for if AI intensity grows and an indicator for no change in AI intensity. The instrument for AI intensity is headline length, which indicates whether newspaper headlines about equity i in quarter t were shorter than the median headline length. Below the coefficient estimates are robust standard errors clustered at the equity level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively. For a detailed description of each variable, see the definitions in Appendix A.

	Dep. Var. =					
	Top 25% of Accuracy & Quits			General Experience		
Panel A. Talent Pool for Analysts	(1)	(2)	(3)	(4)	(5)	(6)
Artifical Intelligence (AI) Intensity	0.075***			-0.078***		
	(0.022)			(0.024)		
Growing AI Intensity		1.027***			-0.938***	
		(0.376)			(0.355)	
No Change in AI Intensity			-2.031**			1.856**
			(0.829)			(0.772)
Additional Controls	Y	Y	Y	Y	Y	Y
Time Fixed Effects	Y	Y	Y	Y	Y	Y
First Stage F-Stat	195.9	35.5	15.7	195.9	35.5	15.7
T-Stat on Instrument	14.0	2.7	2.4	14.0	2.7	2.4
Adjusted R-squared	1%	1%	1%	74%	70%	70%
Observations	81,597	63,544	63,544	81,597	63,544	63,544

	Dep. Var. =					
	Bold Revision			Distance from Consensus		
Panel B. Strategic Response by Analysts	(1)	(2)	(3)	(4)	(5)	(6)
Artifical Intelligence (AI) Intensity	0.015			-0.077***		
	(0.019)			(0.026)		
Growing AI Intensity		0.191			-1.278***	
		(0.255)			(0.413)	
No Change in AI Intensity			-0.377			2.528***
			(0.511)			(0.945)
Additional Controls	Y	Y	Y	Y	Y	Y
Time Fixed Effects	Y	Y	Y	Y	Y	Y
First Stage F-Stat	195.9	35.5	15.7	195.9	35.5	15.7
T-Stat on Instrument	14.0	2.7	2.4	14.0	2.7	2.4
Adjusted R-squared	19%	22%	22%	5%	6%	6%
Observations	81,597	63,544	63,544	81,597	63,544	63,544

Table 7. Market Reaction to Analysts' 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 dowgrade to a sell or strong sell. Returns are in excess of benchmark portfolios matched on size, book-to-market, and momentum. Downgrades have been multipled 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. Panel A displays abnormal returns over the main sample period from 2010Q1 to 2016Q3, whereas Panel B includes recommendations from the prior five years as a further benchmark. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively. For a detailed description of each variable, see the definitions in Appendix A.

	Dep. Var. = Excess Returns		Dep. Var. = Excess Volume	
	[0,1]	[0,5]	[0,1]	[0,5]
Panel A. Reaction to Analyst Recommendations	(1)	(2)	(3)	(4)
Artifical Intelligence (AI) Intensity	-0.24%***	-0.27%***	-0.047***	-0.034*
T-stat	(5.01)	(5.25)	(5.91)	(1.76)
Time Fixed Effects	Y	Y	Y	Y
Analyst Fixed Effects	Y	Y	Y	Y
Observations (Recommendations)	39,454	39,454	39,454	39,454

	Dep. Var. = Excess Returns [0,1]			
Panel B. Change in Reaction Pre & Post FinTech Entry	(1)	(2)	(3)	(4)
Artifical Intelligence (AI) Intensity	-0.28%***	-0.25%*	-0.10%*	-0.05%
T-stat	(3.36)	(1.87)	(1.93)	(0.54)
Time Fixed Effects	Y	Y	Y	Y
Analyst Fixed Effects	Y	-	Y	-
Firm Fixed Effects	Y	-	Y	-
Firm-Analyst Pair Fixed Effects	-	Y	-	Y
Include 5 Years Before FinTech	N	N	Y	Y
Observations (Recommendations)	35,790	18,905	75,165	42,087
Firm-Analyst Pairs	-	7,593	-	14,418

INTERNET APPENDIX

Appendix A. Variable Definitions

We use data from IBES, Zacks, CRSP, Compustat, and Thomson Reuters to construct our financial analyst sample. To construct our various measures of accuracy and bias, we use diluted, U.S. currency quarterly earnings per share (EPS) forecasts from one to eight quarters out as well as diluted, U.S. currency annual EPS forecasts from one to two years out. The remaining EPS forecasts that are greater than two years out or more than eight quarters out represent less than 2% of the universe of forecasts and are not well populated to evaluate the consensus; hence, we excluded them. We include in our set of forecasts those that are original forecasts, announced confirmations of previous forecasts, and revised forecasts. Each variable is winsorized at the 1st and 99th percentile to mitigate the influence of extreme observations. Definitions are as follows:

Mean (Median) Bias as a Percentage of the Absolute Value of Consensus EPS is the difference between the analyst's forecast and the actual EPS divided by the absolute value of the consensus EPS for equity i in quarter t. Because our analysis is conducted at the equity level, we further aggregate forecast biases and consider the consensus bias expressed as the mean (median) bias among all analysts covering a particular equity.

Mean (Median) Accuracy as a Percentage of the Absolute Value of Consensus EPS is the absolute value of the signed forecast error (i.e., the difference between the analyst's forecast and the actual EPS) divided by the absolute value of the consensus EPS for equity i in quarter t. Because our analysis is conducted at the equity level, we further aggregate forecast accuracy and consider the consensus accuracy expressed as the mean (median) forecast error among all analysts covering a particular equity.

Mean (Median) Bias as a Percentage of the Previous Quarter's Stock Price is the difference between the analyst's forecast and the actual EPS divided by the closing price for equity 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. Because our analysis is conducted at the equity level, we further aggregate forecast biases and consider the consensus bias expressed as the mean (median) bias among all analysts covering a particular equity.

Mean (Median) Accuracy as a Percentage of the Previous Quarter's Stock Priceis the absolute value of the signed forecast error (i.e., the difference between the analyst's forecast and the actual EPS) divided by the closing price for equity 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. Because our analysis is conducted at the equity level, we further aggregate forecast accuracy and consider the consensus accuracy expressed as the mean (median) forecast error among all analysts covering a particular equity.

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

Forecast Dispersion is the standard deviation of all analyst forecasts covering stock i in quarter t. $(VALUE)$

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 $= log(\frac{PRCC\_F \times CSHO + DLC + DLTT + PSTKL - TXDITC}{AT})$

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

Volatility of ROE comes from estimating an AR(1) model for each equity'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.

Institutional Ownership data come from Thomson-Reuters via 13F SEC filings. Ownership percentages are based on the number of shares outstanding and correspond to calendar dates.

Hedge Fund Ownership data come from Factset. (IO_CAT6)

Affiliated Analyst is an indicator variable for whether an analyst works at a brokerage house with a pre-existing relationship with the firm through business underwriting an IPO, SEO, or as an advisor on an M&A deal. IPO, SEO, and M&A deal data are pulled from SDC.

Brokerage Size is the number of analysts at the brokerage firm.

Brokerage Prestige is an indicator variable that takes a value of one if the brokerage firm is listed that year as one of Institutional Investor Magazine's top brokerage houses.

Firm 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 Firms 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.

Days Since Last Forecast is the average number of days elapsed since the most recent forecast for that same stock by i by analyst j in a given quarter t.

Forecast Horizon is the average number of days between the estimate date and the reference date, which is the fiscal period end date, in a given quarter t for a stock i covered by analyst j.

Forecast Frequency is the number of forecasts for stock i issued by analyst j during the previous year.

Bold Revision 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.

Distance from Consensus is the average distance from the consensus for forecast revisions by all analysts covering stock i in quarter t. This variable is a continuous measure of the boldness of the forecast revision. To define this continuous measure of boldness, we follow the construction in Clement and Tse (2005). We calculate the distance of analyst j's revised forecast for firm i from the pre-revision consensus forecast. We take the absolute value of the distance of the revised forecast from the consensus minus the minimum absolute distance for analysts who follow firm i in quarter t, with this difference scaled by the range in absolute distances for analysts following firm i in quarter t.

Analyst Quits 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.

Analyst Quits and Is in Top % of Accuracy is an indicator variable for if analyst j stops appearing in the IBES dataset altogether and was in the top 10% (or 25%) of accuracy across all analysts in quarter $t-1$. For each stock i in quarter t, the analyst with the minimum signed forecast error is identified as the most accurate. Then, we calculate for what percentage of stocks the analyst covers for which he or she had the most accurate forecast (e.g., if an analyst covers five stocks and is most accurate one time, the analyst is most accurate 20% of the time). Then, based on analyst-quarter observations for our sample period, we define the percentiles for the percentage of time most accurate.

Analyst Initiates Coverage is an indicator variable for if analyst j (defined as a unique $ANALYS$ in IBES pneumonics) begins covering an equity not previously covered in the IBES database.

Analyst Ends Coverage is an indicator variable for if analyst j (defined as a unique $ANALYS$ in IBES pneumonics) ends covering an equity in the IBES databse. 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.

Analysts Contribution to Price Informativeness is estimated as $AC_{j,t} = \frac{\sum_{d=1}^{NREVS} |Ret_{j,d} - DecRet_{j,d}|}{\sum_{d=1}^{63} |Ret_{j,d} - DecRet_{j,d}|}$ where d denotes trading days in a quarter, $NREVS$ denotes the number of unique days for which there is at least one analyst forecast, j denotes firm, and t denotes quarter. To mitigate the potential concern for the AC measure of analysts "piggybacking" off of management or other experts analysis, we exclude from the numerator days when earnings announcements are made and days when bloggers post about the equity.

Appendix B. Additional Tables

Table B1. Summary Statistics for Newspaper Headlines

This table provides summary statistics related to newspaper headline length by year for the sample period from 2009 to 2016. Headline data comes from Ravenpack. The following newspapers are included in the analysis: USA Today, the Wall Street Journal, the New York Times, the Los Angeles Times, the Chicago Tribune, the Washington Post, the Financial Times, and the DowJones Newswire.

Year (1)	Obs. (2)	Mean (3)	25th Percentile (4)	Median (5)	75th Percentile (6)
2009	619,854	55.4	50	56	60
2010	1,119,096	52.6	45	53	60
2011	988,949	55.4	47	56	62
2012	1,055,145	57.0	48	58	63
2013	846,705	56.1	48	57	63
2014	1,004,774	57.8	49	58	65
2015	998,023	58.6	49	58	66
2016	905,906	58.9	48	58	67
All	7,538,452	56.5	48	57	63

Table B2. Example Newspaper Headlines

This table presents example headlines for Apple that fall at exactly at the 25th percentile and 75th percentile of headline length, respectively.

25th Percentile of Headline Length	75th Percentile of Headline Length
Apple’s Schiller Rides to the App Store’s Rescue	Apple to Unveil New Software Tuesday; Netbook Speculation Swirls
Apple Mac Sales Did Better Than Expected in June	Apple CFO: Consumer Spending Stronger Than Business, Government
Nokia Sues Apple Over iPhone Patent Infringement	Elan to Expand Patent Lawsuit Against Apple to Include the iPad
Analysts Ask If the iPad Can Live Up to Its Hype	Apple Urges Respect for Workers After Death Over Missing iPhone
Apple Adds 2 Publishers to Its Store for E-Books	iPad 3g Shortage Spurs At&T Discussions About Its New Data Plan
Apple Unveils iPad Tablet With Onscreen Keyboard	Apple Launches Ad System for Mobile Devices in Race With Google
Tablet? Slate? New Devices Emerge As Apple Looms	FTC Said to Prepare Review of Apple Tactics in Mobile Ad Market
Apple’s Mac Sales Shine, iPhone Lags Street View	Jobs Wears Tech Crown As Apple Eclipses Microsoft in Market Cap
Apple Still Underdog in China, Despite New Store	Android Sales Top Apple iPhone in First-Half 2010, Says Nielsen
Are Hedge Funds Hoping for an Apple ‘Slingshot’?	Apple Closes In on Nintendo With 40 Million iPod, iPhone Gamers
Apple Gives App Developers Its Review Guidelines	Did Steve Jobs Accidentally Confirm Cameras for Next Gen iPads?
Apple Gearing Up for Newspaper Subscription Plan	Once Again, Apple’s New Design Won’t Accommodate Your Old Cords

Table B3. 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 FE	Y
Adjusted R^2	0.10%
Observations	7,538,452

Table B4. 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.

	Dep. Var. = Headline Length (1)	R-squared 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	