Decision Fatigue and Heuristic Analyst Forecasts*

David Hirshleifer
Paul Merage School of Business
University of California, Irvine

Yaron Levi[†]
Marshall School of Business
University of Southern California

Ben Lourie
Paul Merage School of Business
University of California, Irvine

Siew Hong Teoh
Paul Merage School of Business
University of California, Irvine

July 20, 2018

Abstract

Psychological evidence indicates that decision quality declines after an extensive session of decision-making, a phenomenon known as *decision fatigue*. We study whether decision fatigue affects analysts judgments. Analysts cover multiple firms and often issue several forecasts in a single day. We find that forecast accuracy declines over the course of a day as the number of forecasts the analyst has already issued increases. Also consistent with decision fatigue, we find that the more forecasts an analyst issues, the higher the likelihood the analyst resorts to more heuristic decisions by herding more closely with the consensus forecast, by *self-herding* (i.e., reissuing their own previous outstanding forecasts), and by issuing a rounded forecast. Finally, we find that the stock market understands these effects and discounts for analyst decision fatigue.

JEL classification: G02, G14, G2.

Keywords: Analysts, Decision Fatigue, Heuristic, Behavioral Finance.

---

*We thank David Aboody, Kenneth Ahern, James Choi, Michael Clement, Dan Givoly, Lisa Kramer, Stanimir Markov, John Matsusaka, David Solomon, Brett Trueman and an anonymous referee; seminar participants at UC Irvine, the American Accounting Associations Annual Meeting, the UCLA-USC-UCI Accounting Research Conference, and the Law and Finance Conference at the University of San Diego, Society of Financial Studies Cavalcade North America for helpful comments.

[†]Corresponding author. Email: ylevi@marshall.usc.edu, phone: 213-821-4160.

1. Introduction

The literature on the determinants of analyst forecasting behavior (e.g., Clement 1999; Bradshaw 2011) emphasizes the errors that derive from conflicts of interest (e.g. Kothari et al. 2016; Mola and Guidolin 2009; Ljungqvist et al. 2007; Kirk 2011; Christophe et al. 2010), and psychological bias (see Ramnath et al. 2008 for a review of the literature). We test here whether analyst decision fatigue (i.e., a decline in decision quality after an extensive session of decision-making) affects forecasting behavior. Specifically, we investigate whether the number of forecasts an analyst has already made during a given day affects the accuracy of the next forecast the analyst makes on that same day. We also test whether analysts who have issued more forecasts during a given day behave more heuristically in the form of herding in their forecasts toward the consensus forecast, self-herding (i.e., reissuing their own outstanding forecasts) or providing a rounded forecast (forecast that ends with zero or five).

A large body of evidence in psychology suggests that judgments and decisions made under greater pressure, distraction, or fatigue tend to be made more heuristically. The distinction between heuristic and non-heuristic decision-making can be understood using the classification of judgments and decisions emphasized by Kahneman and Egan (2011) and initially introduced by Stanovich and West (2000). In this model, decisions arise either from System 1, in which the decision is made using quick and easy intuitive cognitive processes, or from System 2, in which decisions are the result of slow, rigorous reasoning processes. System 2 thinking (i.e., non-heuristic decision-making) requires more mental resources, so individuals tend to switch to System 1 thinking (i.e., heuristic decision-making) after an extended period of System 2 thinking.

We expect that analysts who use System 2 thinking will produce higher-quality forecasts than System 1 thinking. We also predict that when analysts become mentally fatigued, they will exhibit a reduced ability to issue an accurate forecast and are more likely to use heuristics (System 1 thinking) when issuing a forecast. These heuristics include techniques

such as conforming to the consensus or reiterating a previous forecast.

Baumeister et al. (1998) describe decision fatigue as a consequence of “ego depletion,” defined as a draining of mental resources. They argue that the self-control required for careful cognitive processing and systematic decision-making requires mental resources that are in limited supply. Self-control is typically impaired when the cognitive resources available for decision-making are low. Thus, when people devote effort to complex decisions over a given period of time, the resulting decision fatigue temporarily reduces the quality of their subsequent decisions. Many subsequent studies in psychology have provided further evidence in support of decision fatigue (Baumeister and Tierney 2012).

There is also anecdotal evidence that professionals are aware and concerned enough about the negative effects of decision fatigue to take active steps to counteract it. For example, President Barack Obama has explained that he minimizes his food and clothing choices to improve his other decisions (Lewis 2012). Steve Jobs and Mark Zuckerberg famously wear only limited styles and colors of clothing. Managers at hedge fund Voss Capital wrote to investors that they encourage their employees to take frequent breaks, even intraday naps or meditation, to prevent overuse of System 1 (“thinking fast”) and to avoid making mistakes (Wadhwa 2016).

However, the recent controversy over the reproducibility of experimental studies in social psychology and other fields has also engulfed the large experimental literature on ego depletion. For example, a large-scale multi-lab experimental study finds no discernible ego depletion effects (Hagger et al. 2016).¹ Nevertheless, even Hagger and Chatzisarantis (2016) in their rejoinder state that, “For the record, we think that ego-depletion is a ‘real’ phenomenon analogous to cognitive fatigue.” They conclude with a call for further study of the topic. One of the contributions of our study is to evaluate decision fatigue using archival data on analyst forecasts instead of laboratory experiments.

---

¹ This triggered commentaries on the study by Baumeister and Vohs (2016) and Sripada et al. (2016) criticizing the strength of the treatments used to induce fatigue and raising other statistical issues in the Hagger et al. (2016).

Our study is not the only one to use archival data to test for decision fatigue. The effects of decision fatigue on decision-making have also been documented in a wide variety of settings in other literatures such as political science (e.g., voting in Augenblick and Nicholson 2015) and, more recently, in economics (e.g., purchasing a car in Levav et al. 2010). Decision fatigue has also been shown to be important in major life-changing decisions; for example, Danziger et al. (2011) report that parole judges rule less favorably toward prisoners as the morning approaches lunchtime and as the afternoon approaches the end of the workday.

However, evidence as to whether decision fatigue affects professionals in the capital market setting is very limited. Hirshleifer et al. (2009) provide evidence that on days when relatively more firms announce earnings, there is a stronger post-earnings announcement drift which can be interpreted as consistent with investor decision fatigue. However, the authors interpret this finding as a result of limited attention. Our goal is to test specifically for decision fatigue effects in a professional capital market setting.

For several reasons, analyst earnings forecasting provides an attractive context for studying decision fatigue. First, analysts' errors can be directly measured, allowing us to test for degradation of decision quality. Second, analysts often make forecasts of multiple firms in a single day, so it is feasible to test how the forecasting behavior of an analyst varies with the number of forecasts she has already issued that day. Therefore, the number of recently issued forecasts provides a proxy for analyst decision fatigue. Third, firms are often followed by several analysts. This allows us to measure forecast accuracy for the analyst relative to the consensus forecast. By using a measure of relative forecast accuracy, we can mitigate firm characteristic effects on forecast accuracy to isolate decision fatigue effects more successfully.

During the period for which data on the time issuance of individual analyst forecasts are verified (i.e., 2002-2015), we find strong evidence that is consistent with the negative effects of analyst decision fatigue on the accuracy of one-year-ahead EPS forecasts. Forecasts by analysts are less accurate when they are issued after the analysts have issued a greater number of forecasts for other firms that day.

We further investigate whether forecasts are made more heuristically as the analyst issues more forecasts that day. We find that forecasts of decision-fatigued analysts exhibit greater herding toward the prior consensus forecast. There is also greater self-herding, which means that the forecasts are also more likely to be reissuances of the analyst’s own previous forecast of a firm. We also find that decision-fatigued analysts are more likely to issue a rounded forecast (one whose last digit is zero or five), although owing to sample size, this finding is less robust. These results are consistent with fatigued analysts switching to System 1 thinking, i.e., decisions that are more non-reflective.

Finally, we study whether investors understand and discount for the lower accuracy of forecasts issued when analysts are more fatigued. We do this by testing how the sensitivity of cumulative abnormal returns to forecast revisions by the analyst varies with the number of forecasts of other firms the analyst has already issued that day. We find that the market understands the potential effect of decision fatigue on analyst forecasts: The market reacts less strongly to analysts’ forecast revisions that are made when the analysts are decision-fatigued.

In our tests of decision fatigue of analysts, we assume that fatigue increases with the number of forecasts the analysts have already issued that day, and that the forecasts are issued in the same order that they are being worked on during the day. The first assumption is highly intuitive. The second is consistent with past evidence that suggests that analysts work in a highly time-sensitive environment, which would pressure analysts to issue forecasts as soon as they are finalized (O’Brien and Bhushan 1990; Altinkılıç et al. 2013; Groysberg and Healy 2013).

It is possible that a forecast issued after other forecasts have been issued that same day may actually have been developed earlier in the day (or week) before the analyst became fatigued, or by other non-fatigued analyst team members. Any such time lags between an analyst’s work and the issuance of her forecast biases against obtaining non-null findings (i.e., that decision fatigue has no effect).

Yet another alternative explanation, one that we cannot rule out entirely, is that analysts choose to structure their workday by first working on forecasts for which they have high-quality information relative to the consensus. This would explain both the higher accuracy of early forecasts and the lower tendency in such forecasts toward herding or self-herding. However, it is not obvious why analysts would follow such a work strategy. It may make sense for an analyst to prioritize making forecasts for firms for which the analyst has better information. However, this could just as easily entail making a well-informed forecast at the end of a workday or deferring the ill-informed forecast for the start of the next workday. Nevertheless, to mitigate this concern, in our robustness tests we remove all forecasts that follow an earnings announcement, and we find that our results are similar, both qualitatively and quantitatively. This suggests that our results are not driven by new public information about firms that is not embedded in the consensus.

As contrasted with decision fatigue, physical fatigue might also influence the ability or willingness of an analyst to exert effort and produce an accurate forecast. We address this possibility by including the time of day as a control variable to proxy for physical fatigue. Our main results also hold in subsamples that include only forecasts that were issued before or after noon.

This paper draws from the literature on decision fatigue (e.g., Levav et al. 2010) and analyst forecast accuracy and herding (e.g., Clement and Tse 2005) to examine whether and how decision fatigue affects analyst forecast behavior and to examine the resulting stock market implications (e.g., Givoly and Lakonishok 1979). Our study contributes to three strands of literature. The first strand is the scant literature on decision fatigue in professional settings, which we expand by showing that information intermediaries are affected by decision fatigue. Second, we contribute to the literature on analyst forecast accuracy and herding by showing that analyst forecasting behavior is influenced by the number of forecasts she issued during the same day. Third, we provide evidence about market efficiency by documenting that the market understands the effect of decision fatigue on analyst forecasts.

2. Hypotheses

Extensive evidence from psychology indicates that judgments and decisions that are made under greater pressure, distraction, or fatigue tend to be made more heuristically. This can be described in the terminology of Kahneman and Egan (2011) as greater use of System 1 thinking. Baumeister et al. (1998) propose that willpower is required to maintain attentional focus for decision-making and, like muscle strength, willpower is temporarily depleted by use. Self-control and judgment are impaired when available psychic resources are low.

Several papers have documented the effects of decision fatigue on decision-making. In four laboratory studies, Vohs et al. (2008) find that participants who made choices among consumer goods or college course options suffered from reduced self-control (i.e., less physical stamina, reduced persistence in the face of failure, more procrastination, and lower quality and quantity of arithmetic calculations). However, others who thought about these same options without making choices did not suffer this reduction in self-control. Augenblick and Nicholson (2015) conducted a field study and found that voters who face more decisions before a given vote are significantly more likely to abstain or to rely on decision shortcuts, such as voting for the status quo or voting for the candidates who are listed first on the ballot. Similarly, Levav et al. (2010) show that consumers who are purchasing a car are more likely to choose default levels of attributes when they begin with attributes that offer a greater number of configuration options than when they begin with attributes that offer a smaller number of options.

The psychology literature on ego depletion describes a limited mental resource (ego) analogous to energy or strength, which is consumed by actions that require application of self-control and/or decision-making. Effort consumes this limited resource while using System 2 for decision-making.

A further question that arises in the literature is whether the observed decline in performance after a session of decision-making (decision fatigue) stems from the exhaustion of

the limited resource or from a need to preserve the remaining stock of the resource.² If decision fatigue stems from the exhaustion of the resource, then effort and decision fatigue are exclusive concepts, and a decision fatigued analyst would not be able to exert more effort to improve the quality of his work. If, however, decision fatigue stems from the need to preserve a limited resource, then decision fatigue and effort are not exclusive concepts, and a decision-fatigued analyst might be able to exert more effort and improve the accuracy of his forecast if properly incentivized. Our paper is not able to distinguish between these two possible sources of decision fatigue.

To our knowledge, the only paper to examine the effect of decision fatigue in a professional setting is that of Danziger et al. (2011). The authors studied the proportion of parole requests approved by eight parole judges in Israel in relation to the time since their last meal break. This proportion spikes after each meal, when about 65% of requests are granted (relative to an average of 35%). During the roughly two hours before the judges' next meal, their approval rate drops steadily to about zero just before the meal. It seems that tired and hungry judges tend to fall back on the easier default position of denying requests for parole. This evidence does not, however, distinguish among decision fatigue, physical fatigue, and hunger as sources of heuristic decision-making.

Previous studies of decision fatigue examine nonprofessionals who face a task either for the first time or infrequently, and their level of motivation to perform well in the task successfully is debatable. In contrast, professionals are experts in the task at hand and are highly motivated to perform well. There are reasons to believe ex-ante that professionals might be less subject to decision fatigue. First, they might develop the mental resources required to perform the task to a degree that the fatigue would not play a role (like building a muscle; see Muraven and Baumeister 2000). Second, they might develop a system of heuristics that leads to the right decision ("a sixth sense", as described by Kahneman and Egan 2011). Third, they or their firms might be aware of their decision fatigue and implement a system of checks or

---

² See the survey of the ego exhaustion or preservation literature in Evans et al. (2016).

supports to ensure that decision quality does not drop. So whether professionals are subject to decision fatigue is an interesting empirical question.

The equity analyst setting has some distinctive features that are especially well-suited for testing the effects of decision fatigue. Different analysts will issue different numbers of forecasts earlier in a given day; therefore, the presence of other analysts who cover the same firm at the same time offers a counterfactual benchmark to the forecast being evaluated. Unlike most professions, the outcome of an analyst's decision can be reliably measured: We can observe *ex post* how close the forecast was to the actual result. Finally, analysts work in a highly time-sensitive environment. As such, it is likely that most of their work is performed sequentially, forecast by forecast, with work on any given forecast closely followed by the issuance of that forecast.

Several findings provide support for this interpretation. O'Brien and Bhushan (1990) describe a customer-supplier relationship between financial institutions and brokerage houses. To the extent that institutional investors demand timely information to make trading decisions, financial analysts have incentives to provide prompt forecast revisions to financial institution clients. The evidence in Altinkılıç et al. (2013) suggests that analysts release recommendations soon after the release of new information that has a material effect on the stock price. Groysberg and Healy (2013) report that analysts issue on average 12 notes to every one report, and each note only requires a few hours to write. It is important to note that even if the research is conducted by teams (i.e., associate analysts take part in the process of analyzing a company), a bottleneck is still created by the senior analyst, who signs off on the report and is responsible for communicating the report to the public. When the senior analyst is fatigued and unable to invest the necessary mental resources to reviewing the work done by the team, the senior analyst might resort to more heuristic behavior.

Clement (1999) shows that factors such as analysts' ability, available resources, and portfolio complexity significantly influence forecast accuracy. For example, the author shows that forecast accuracy increases with experience (a proxy for ability) and with employer

size (a proxy for available resources), and that accuracy decreases with the number of firms followed (a proxy for portfolio complexity). We contribute to this literature by testing how mental resources affect forecast accuracy, controlling for past known determinants wherever possible. We predict that, with each additional forecast in the sequence, the analyst becomes more fatigued. This fatigue causes the analyst to rely more on System 1 thinking than System 2 thinking when making a decision that reduces forecast accuracy. Accordingly, we hypothesize:

H1: An analyst's relative forecast accuracy decreases with the number of forecasts the analyst has made earlier in the day.

An analyst who is fatigued can resort to some natural heuristic procedures for generating a forecast. One is to herd by issuing a forecast that is close to the consensus forecast. This is a reasonable shortcut to follow when the analyst lacks the cognitive resources to generate much incremental information relative to the consensus. This hypothesis is new to the herding literature, which has focused primarily on information transmission or agency problems.³ We build on this literature by testing how analysts' mental resources are also a determinant of herding behavior. This leads to our second hypothesis:

H2: The likelihood that an analyst herds increases with the number of forecasts the analyst has made during the day.

Another possible heuristic is to stick closely to the analyst's previous outstanding forecast about the firm. When decision fatigue prevents an analyst from generating much useful new information, another reasonable shortcut is to rely more heavily on previous analyses. In the extreme case, the analyst would self-herd by not updating the previous forecast at all. This leads to our third hypothesis:

H3: The likelihood that an analyst reissues an outstanding previous forecast increases

---

³ Welch (2000) documents herding behavior among analysts. Hong et al. (2000) show that herding is economically rational given analysts' career concerns: being wrong when everyone else is wrong is preferable to being wrong when others are correct. Clement and Tse (2005) find that analyst characteristics, especially those that reflect analyst forecast abilities, affect herding behavior.

with the number of forecasts the analyst has made during the day.

Herrmann and Thomas (2005) provide evidence that rounded forecasts (forecasts ending with zero or five) are less accurate than other forecasts. Dechow and You (2012) show that rounded forecasts are more likely to indicate that the analyst is exerting less effort. Decision-fatigued analysts may be tempted to provide a rounded forecast instead of exerting the mental effort required to provide a more accurate estimate. This leads to our fourth hypothesis:

H4: The likelihood that an analyst issues a rounded forecast increases with the number of forecasts the analyst has made during the day.

Past research indicates that sell-side analysts' forecast revisions are important for investor expectations about firms' earnings and for making investment decisions (e.g., Hodge 2003). This conclusion is supported by the substantial average stock market reaction to the release of forecast revisions (e.g., Brown et al. 1985; Gonedes et al. 1976; Givoly and Lakonishok 1979). Furthermore, there is evidence that market reactions to forecast revisions take into account past forecast accuracy and other correlates of current forecast accuracy (e.g., Bonner et al. 2003; Clement and Tse 2005; Gleason and Lee 2003; Michaely and Womack 1999).

This literature suggests that if the market is efficient, it will take into account the effects of decision fatigue on analyst forecast accuracy. For example, investors may directly take into account the number of previous forecasts the analyst has issued during the day. Alternatively, sophisticated investors understand that analyst herding occurs, and they take this into account when evaluating forecasts. Decision-fatigued analysts are more likely to offer forecasts that are similar to the consensus, and this may also lead to more discounting of their forecasts. This leads to our fifth and final hypothesis:

H5: The more forecasts an analyst has issued earlier in the same day, the weaker the reaction of investors when the analyst issues a forecast revision.

We discussed earlier the possibility that analysts intentionally issue their most well-

informed forecasts early in the workday. To consider this in more depth, suppose that especially precise information signals arrive uniformly throughout the workday, and that the analyst tends to work on a firm forecast whenever a precise signal about that firm first arrives. If so, then precise forecasts will be distributed evenly throughout the workday. Now, if we instead suppose that precise signals arrive only toward the end of the workday (perhaps because these signals are the product of analyst effort during the day), then the most precise forecasts will tend to be issued late in the day. This would bias against finding the results we document, and it strengthens our inference from our evidence that decision fatigue is a factor.

Another possible concern arises in a scenario in which analysts generate their most precise signals overnight or over the weekend, and therefore they issue their most precise forecasts at the start of the workday. We cannot rule out this possibility because it would generate results similar to the implications of decision fatigue. For example, firms often make voluntary disclosures outside of trading hours. This encourages the revision of forecasts, which may occur at the start of the next day. However, this is by no means always the case. Zhang (1998) documents that around half of analysts revise their forecasts in the three days following an earnings announcement, so it is clear that such revisions often occur later than the morning of the first day after the announcement. Nevertheless, we perform robustness checks to ensure that our results are not driven by morning revisions the day after earnings announcements.

Also, although morning forecast revisions are issued in response to news received overnight or during the previous weekend, forecasts issued during the rest of the day are also in response to the arrival of public information. So, although this news is informative, the fact that morning forecast revisions make use of new information does not imply that morning forecasts are more informative than forecasts made at other times of the day.

3. Data and Descriptive Statistics

Data on analysts' EPS forecasts were collected from the Institutional Brokers' Estimate System (I/B/E/S) database over the period 2002-2015.⁴ The starting year of 2002 was chosen because this was the first year that the announcement date of the forecast was verified (Hoechle et al. 2015).⁵ Similar to prior literature (e.g., Gleason and Lee 2003; Clement and Tse 2005; Kumar 2010), we focus on one-year-ahead earnings forecasts.

The focus of this paper is on timely forecasts that are issued during the workday. So, we focus on forecasts that were prepared or at least partially prepared during a single day and were released on that day in sequence. Accordingly, we limit our sample to days when the analyst only issued forecasts between the working hours of 9:00 a.m. and 7:00 p.m.⁶ Each forecast issued during the day is marked as a decision by the order in which it was issued.

Table 1 shows the number of analyst's decisions in our sample and the partition between the number of forecasts in a day. On average, analysts make 1.3 forecasts per day (on days when forecasts are issued), and our sample consists of 386,924 total forecasts. On most of the analyst-days in the sample (255,613), the analyst only made one forecast. On 27,975 analyst-days, the analyst made two forecasts, resulting in 55,950 forecasts; the number of analyst-days that have a larger number of forecasts continues to decrease with the number of forecasts.

Our main dependent variables of interest are *ACCURACY*, *HERDING*, *REISSUE*, and *ROUNDING*. Following prior research, we compare the accuracy of an analyst's one-year-ahead EPS forecasts for a particular company at a given time to the mean level of accuracy for all analysts who make forecasts for the same company and time period within

---

⁴ Following previous literature, we exclude utilities and financial services firms (SIC codes 4900-4999 and 6000-6999).

⁵ Bradley et al. (2014) studies the timestamp error of analysts' recommendation revisions relative to hand-collected timestamps provided by newswire sources between 2002 and 2007. The authors estimate that the mean delay in IBES timestamps is 2.38 hours, and the standard deviation is 3.51 hours. Our results hold in the subsample of forecasts with afternoon timestamps, meaning that even with the estimated time stamp error the forecasts in this subsample are not likely to have been issued before the start of the workday.

⁶ Changing the length of the workday provides qualitatively similar results.

a comparable forecast horizon (Jacob et al. 1999; Clement 1999; Hong et al. 2000; Cowen et al. 2006). This controls for any firm- or time-specific factors that affect forecast accuracy. We therefore define

$$RELATIVE\ ACCURACY_{i,j,t} = \frac{Median\ Forecast\ Error\ of\ All\ Analysts_{j,t} - Analyst's\ Forecast\ Error_{i,j,t}}{Standard\ Deviation(Forecast\ Error\ of\ All\ Analysts_{j,t})}$$

where *Analyst's Forecast Error*_{*i,j,t*} is the absolute value of actual earnings minus the earnings forecast of analyst *i* at firm *j* at time *t*, and the *Median Forecast Error of All Analysts*_{*j,t*} is the median EPS forecast error for all analysts who cover firm *j* within the same 90 days. The denominator standardizes across firms by dividing by the standard deviation of EPS forecast errors across all analysts who cover firm *j* at time *t*.

Following Clement and Tse (2005), we define *HERDING*_{*i,j,t*} as a binary variable that receives the value of 1 if analyst *i*'s forecast of company *j* at time *t* is between the consensus forecast at time *t* and the analyst's own previous forecast, and 0 otherwise. (All other variables are defined in Appendix A.)

We also estimate a new measure, *REISSUE*_{*i,j,t*}, which is a dummy variable that takes the value of 1 if a forecast is reissued (self-herding), and 0 otherwise. When an analyst reissues a forecast, I/B/E/S does not create a new record in its dataset. Instead, I/B/E/S collects information on the date (REVDATS) and time (REVTIMS) the analyst reissued the outstanding forecast.⁷ We use this date and time to ascertain when a forecast was reissued.

Following Dechow and You (2012), we define *ROUNDING*_{*i,j,t*} as a dummy variable that takes the value of 1 if a forecast ends with zero or five in the penny digit, and 0 otherwise.

Table 2 shows descriptive statistics by number of forecasts made by the analyst for the given day. As expected, *ACCURACY* declines and *HERDING* increases as the analyst

---

⁷ If an analyst's report does not contain a revision to the forecast, then I/B/E/S does not keep that forecast as a separate record. It retains the original record for that forecast, but updates the review date (REVDATS) and time (REVTIME) for the forecast to make it current. If the forecast is changed, only then does I/B/E/S enter a new record in its database but with a new announcement date (ANNDATS).

makes more forecasts throughout the day. *ROUNDING* is relatively stable and has no clear pattern. The size of the brokerage house and the forecast age (the number of firms the analyst follows) are decreasing (increasing) with each sequential decision. The analyst's experience with the firm and the level of effort invested in a firm do not seem to follow any specific pattern. The type of firm seems to be related to the decision order as well. Firms that are forecasted earlier in the day tend to be smaller, followed by fewer analysts, have lower ROA, higher sales growth, higher R&D, higher fraction of intangible assets, and a lower earnings-to-price ratio.

One possible reason for this phenomenon is that analysts try to first issue forecasts for firms for which the forecasting problem is more complex, or for which the information environment is sparser. This may be valuable to investors who want to trade during the day and who will have the most trouble evaluating such firms until the analyst provides a new and timelier forecast. Alternatively, an analyst may recognize that she will be fatigued later in the day, and therefore try to complete the most challenging tasks much earlier in the day when she is not fatigued.⁸

4. Results

4.1. Accuracy

To assess whether analysts' forecast accuracy decreases as a function of the number of earlier forecasts they have made during the day, we estimate the following regression model:

$$RELATIVE\ ACCURACY_{i,j,t} = \alpha + \beta_1 DECISION\ RANK_{i,j,t} + \beta_2 CONTROLS + \epsilon_{i,j,t} \quad (1)$$

---

⁸ If this is occurring, and to the extent that our controls for determinants of decision accuracy are imperfect, it would tend to cause us to find that earlier forecasts are less accurate than later forecasts. It would therefore bias against the results that we actually find.

Where our key independent variable $DECISION\ RANK_{i,j,t}$ is the logarithm of the number of forecasts an analyst has issued before the focal forecast plus one.⁹ Our controls for other determinants of analysts' relative accuracy include the number of companies covered by the analyst, the brokerage house size, the analyst's firm-specific experience, the age of the forecast, the forecast frequency, and the number of analysts who cover the firm. Finally, we control for the time of day being a measure of physical fatigue rather than decision fatigue.

To test our hypotheses, we estimate Model 1 using three different specifications. The first specification excludes analyst fixed effects. It estimates whether the accuracy of a forecast deteriorates, on average, as a function of the number of forecasts an analyst has previously issued during the day under the implicit assumption that analyst accuracy is ex ante identical across analysts. The second model includes analyst fixed effects to control for analyst differences in accuracy. Thus, the model examines whether on average for a given analyst, the accuracy of the forecast deteriorates as a function of the number of forecasts the analyst has previously issued during the day. Finally, we include in the model analyst-day fixed effects, which compare whether for a given analyst-day the accuracy of the forecast deteriorates as a function of the number of forecasts the analyst has previously issued during that day, which controls for the fact that accuracy may be greater on some days than on others.

The results presented in Table 3, Columns 1 and 2, indicate that on average the accuracy of the forecast deteriorates as a function of the number of forecasts the analyst has previously issued during the day. In Column 2, the coefficient on our key independent variable, $DECISION\ RANK$, is -0.225 and is significant at the 1% level. This suggests that on average a one-unit increase in $DECISION\ RANK$ leads to a forecast that is 0.225 standard deviations less accurate relative to the consensus. This is an economically meaningful effect.

Columns 3 and 4 indicate that, for a given analyst, the accuracy of the forecast deteriorates as a function of the number of forecasts the analyst has previously issued during

---

⁹We winsorize the variable at 5. Results are robust to not winsorizing.

the day. In Column 4, the coefficient on our variable of interest, *DECISION RANK*, is -0.169 and is significant at the 1% level. This suggests that on average a one-unit increase in *DECISION RANK* leads to a forecast that is 0.169 standard deviations less accurate relative to the consensus for the same analyst, regardless of what day or for what type of firm the forecasts were issued.

H1 is formally tested in Columns 5 and 6. By adding analyst-day fixed effects to the regression specification, we test whether for a given analyst-day the accuracy of the forecast deteriorates as a function of the number of forecasts the analyst has previously issued during that day. In Column 6, the coefficient on our variable of interest, *DECISION RANK*, is -0.067, significant at the 5% level. This implies that on average a one-unit increase in *DECISION RANK* leads to a forecast that is 0.067 standard deviations less accurate relative to the consensus for the same analyst and the same day. A different way to interpret the economic magnitude is by examining Table 2. The average accuracy decreases from the first forecast to the second forecast by 0.089, which is equivalent to a decrease of 18.5%.

The results in this section suggest that forecast N is on average more accurate than forecast $N + 1$ and suggests that the quality of decisions deteriorates as a function of the number of previous decisions the analyst has made during that day. This is true for our three test specifications. First, analyst i 's forecast N is more accurate than analyst j 's forecast $N + 1$. Second, analyst i 's forecast N is more accurate than the forecast $N + 1$ on a different day. Third and most importantly, analyst i 's forecast N is more accurate than forecast $N + 1$, which was issued on the same day. All of these comparisons hold constant across firms, and therefore the results are independent of firm characteristics.

4.2. Herding

We now turn to the question of whether analysts who are more decision fatigued resort more to heuristic decision-making. We therefore test whether analysts are more likely to issue a herding forecast as a function of the number of forecasts the analyst has previously

issued during the day. We use the following logistic regression:

$$Pr(HERDING_{i,j,t}) = f(\alpha + \beta_1 DECISION\ RANK_{i,j,t} + \beta_2 CONTROLS + \epsilon_{i,j,t}) \quad (2)$$

Table 4 summarizes the relationship between herding and decision fatigue. We present two regression specifications (logit and fixed effects logit). Both specifications include our set of controls from Model 1. Columns 1 to 6 indicate that an analyst’s issuance of a herding forecast is positively associated with the number of earlier same-day forecasts made by the analyst. This is true for all analysts on average (Columns 1 and 2) and for an analyst who covers a specific firm (Columns 3 and 4).

To formally test H2, we use the conditional form of the logit regression and control for analyst-day FE.¹⁰ The results are presented in Columns 5 and 6. Consistent with the hypothesis, within a specific analyst-day, the analyst is more likely to herd with each sequential decision. The coefficient is equal to 0.086, significant at the 5% level. The marginal effect at the mean is 0.02, meaning that a one standard deviation increase in *DECISION RANK* corresponds to a 0.7% increase in the probability of herding. A different way to interpret the economic magnitude is by examining Table 2. The probability of herding increases from the first forecast to the second forecast by 0.023, which is equivalent to an increase of 8.27%.

4.3. *Reissued forecasts*

Another possible heuristic is that the analyst would self-herd by not updating the previous forecast at all. We therefore test whether analysts are more likely to reissue an outstanding forecast as a function of the number of forecasts the analyst has previously issued during the day. We use the following logistic regression:

---

¹⁰ We use conditional logit in order to estimate the fixed effects model consistently. By conditioning the likelihood on the number of successes in each panel, we avoid estimating the coefficients of the fixed effects themselves. As a result, this procedure produces consistent estimates of the remaining coefficients.

$$Pr(REISSUE_{i,j,t}) = f(\alpha + \beta_1 DECISION\ RANK_{i,j,t} + \beta_2 CONTROLS + \epsilon_{i,j,t}) \quad (3)$$

The results are reported in Table 5. Consistent with the hypothesis, the coefficient of *DECISION RANK* is positive and significant across all specifications. H4 is formally tested in Columns 5 and 6. The results suggest that within a specific analyst-day, the more forecasts the analyst has issued previously during the same day, the more likely the analyst is to self-herd by reissuing an outstanding previous forecast. The marginal effect at the mean is 0.262, meaning that a one standard deviation increase in *DECISION RANK* corresponds to an 8.2% increase in probability of reissuing the same forecast within a given analyst day. A different way to interpret the economic magnitude is by examining Table 2. The results show that the probability of reissuing the same forecast increases from the first forecast to the second forecast by 0.064, which is equivalent to an increase of 11.1%.

4.4. *Rounding*

As decision fatigue increases, the analyst might be tempted to be heuristic in other ways as well. One example is providing a rounded estimate instead of devoting the additional mental effort required for a more accurate estimate. We therefore test whether analysts are more likely to provide a rounded forecast as a function of the number of forecasts the analyst has previously issued during the day. We use the logistic regression:

$$Pr(ROUNDING_{i,j,t}) = f(\alpha + \beta_1 DECISION\ RANK_{i,j,t} + \beta_2 CONTROLS + \epsilon_{i,j,t}) \quad (4)$$

We split the sample into earning forecast below \$1, and forecasts in the range of \$1 to \$10. (There are less than 60 cases with earnings forecasts above \$10.) Consistent with Dechow and You (2012), we find no effect when the level of earnings forecast is below \$1, probably since rounding the penny digit would lead to considerable forecast inaccuracy for

such stocks. In other words, even an analyst who is behaving heuristically is unlikely to make such an extremely crude estimate.

The remaining sample of earning forecasts between \$1 and \$10 is considerably smaller than for our main sample, reducing statistical power. The results for this sample are in Table 6. The coefficient of *DECISION RANK* is generally positive and increases in both magnitude and significance as we improve the model specification to include controls and analyst-day fixed effects. H4 is formally tested in Column 6.

These results suggest that within a specific analyst-day, the more forecasts the analyst has issued previously during the same day, the more likely the analyst is to provide a rounded forecast. The marginal effect at the mean is 0.033, so that a one standard deviation increase in *DECISION RANK* corresponds to a 0.7% increase in probability of rounding the forecast in a given analyst-day. The economic magnitude can also be estimated using Table 2, which shows that the probability of rounding a forecast increases from the first forecast to the second forecast by 0.003, an increase of 1%.

4.5. Market Reaction

To examine whether investors react differently to forecast revisions issued by analysts as a function of the number of earlier forecasts they have made during the day, we estimate the following regression:

$$\begin{aligned}
 CAR_{i,j,t} = & \alpha + \beta_1 DECISION\ RANK_{i,j,t} + \beta_2 FORECAST\ REVISION_{i,j,t} \\
 & + \beta_3 DECISION\ RANK_{i,j,t} * FORECAST\ REVISION_{i,j,t} + \beta_4 CONTROLS \\
 & + \epsilon_{i,j,t}
 \end{aligned} \tag{5}$$

where $CAR_{i,j,t}$ is the 3-day market-adjusted excess return for firm j centered on the forecast revision issued by analyst i at time t . The variable $FORECAST\ REVISION_{i,j,t}$ is a measure of the difference between the current annual earnings forecast for analyst i

following firm j at time t and the annual earnings forecast issued immediately before the current annual earnings forecast, scaled by the standard deviation of forecasts of all analysts who cover firm j at time t .

To calculate the forecast revision, we require that the analyst issue both a current and a prior annual earnings forecast for the same firm and year. We choose the analyst's prior forecast to calculate the forecast revision, because it is more informative to the market than the consensus forecast (Gleason and Lee 2003; Stickel 1991).

Table 7 reports the results from estimating Model 5. As expected, the estimated coefficient on *FORECAST REVISION* is positive and statistically significant regardless of the specification used, which indicates that the market reaction to the release of the revised forecast is associated with the signed magnitude of the forecast revision. Consistent with Hypothesis 3, the estimated coefficient on *FORECAST REVISIONDECISION RANK* is negative and statistically significant in all specifications. The coefficient on our variable of interest ranges from -0.007 without control variables and fixed effects to -0.001 when including all control variables and analyst-day fixed effects. The economic significance seems large; for example, in Column 5, the coefficient of *FORECAST REVISIONDECISION RANK* is -0.002, and it is equal to 20% of the coefficient of *FORECAST REVISION*.

This finding indicates that the market reacts less strongly the more prior forecasts the analyst has made during the day. Previous research shows that investors respond less aggressively to forecast revisions that derive from herding behavior (Gleason and Lee 2003). Investors also underweight rounded forecasts in forming expectations, indicating that investors regard these forecasts as less reliable (Dechow and You 2012). Since we find that higher-ranked forecasts are more likely to reflect herding and rounding behavior, our finding that the market response to higher rank forecasts is weaker is consistent with this previous research.^{11, 12}

---

¹¹ There are rational settings in which herding or cascading is a rational means of exploiting the information possessed by earlier decision-makers (e.g, Banerjee (1992); Bikhchandani et al. (1992)). However, irrational herding as induced, for example, by decision fatigue, will tend to reduce the quality of decisions.

¹² We also test whether investors pay attention to the forecast rank itself or only to herding and rounding

4.6. *Alternative Explanations, Robustness, and Additional Tests*

Our tests are based on the ordering of forecasts exogenously inducing variation in decision fatigue. However, the ordering might not be random; analysts may rank firms and prioritize some over others for earlier forecasts. This might induce variation in the accuracy of forecasts in relation to rank for reasons other than decision fatigue. Of course, a non-random ranking rule is not in itself inconsistent with decision fatigue. But we now consider whether such prioritization affects the inferences from our tests.

We explore two types of potential ranking rules:

1. A ranking rule with a constant firm ordering, meaning that some specific firm A always ranked ahead of firm B for forecasts by a specific analyst. In other words, if information about the two firms arrives at about the same time, the analyst issues a forecast for firm A before firm B. Examples for such ranking rules include starting with firms which are more (or less) complex, firms that have relationships with the analysts' employer, firms that are more important to investors, or firms in which the analyst specializes.

For the sake of specificity, we consider first ranking more important firms higher. Crucially, analysts issuing forecasts of more important firms *earlier in time* does not imply that they will issue such forecasts *earlier in the day*. To see this, suppose, for example, that there was no decision fatigue, and that as new information arrived, analysts issued timelier forecasts for more important firms. This would in some cases push a forecast earlier in the day, if news arrived early in the day. But in other cases, it would push forecasts to late in the day if news arrived late in the day, instead of deferring to early the next day. Overall, prioritization of more important firms has no obvious implication for whether important firms on average receive updates earlier or

---

behaviors by repeating the tests in Table 7 with the variables *ROUNDING*, *HERDING*, and their interactions with *FORECAST REVISION*. Consistent with Gleason and Lee (2003), we find that the market response is weaker for forecast revisions that represent herding behavior. *ROUNDING* is not significant in this analysis. The coefficients on the interaction of *DECISION RANK* and *FORECAST REVISION* are all negative and significant, and the magnitudes are identical to the coefficients reported in Table 7 when including analyst-firm or analyst-day. This evidence suggests that investors pay attention to the rank of the forecast itself in addition to their response to the usage of heuristics by the analyst.

later in the day.

Consider now the case in which there is decision fatigue, and analysts understand this. In this case, there is a reason for analysts to push more important firms to earlier in the day (perhaps even by deferring the update for an important firm from late in the day to early the next day). However, this argument is premised on the idea that there is decision fatigue. So, this argument does not seem to provide a counterargument to the conclusion of the paper that there is decision fatigue.

Nevertheless, to verify robustness, we explicitly tested if analysts follow a constant decision rule and found close to zero autocorrelation coefficient for a given analyst-firm pair.¹³

Also, our relative accuracy measure, which compares the accuracy for a particular company and time period to the median accuracy for all analysts who make forecasts for the same company and time period within a comparable forecast horizon, mitigates any firm characteristic effects (Jacob et al. 1999; Clement 1999; Hong et al. 2000; and Cowen et al. 2006). Nevertheless, for robustness, we repeat our main analysis with analyst-firm-year fixed effects. Results are summarized in Table ???. The results for the effects of decision fatigue on the accuracy, herding, and reissuance are robust to the inclusion of analyst-firm-year fixed effects. The coefficient of *DECISION RANK* in the rounding regression (column 4) is positive but not significant, but this is not surprising given the smaller sample size and resulting lower power of this test. So, our results are robust to non-random ranking caused by firm and analyst characteristics.

2. A ranking rule with a non-constant firm ordering. In such a rule, an analyst follows some state variable which determines whether firm A is ranked before firm B. For example, an analyst might order his forecast based on when new information about each firm arrives.

---

¹³ We tested the autocorrelation coefficient using a sample which includes only days in which an analyst made two forecasts. We regress *DECISION RANK* on lagged *DECISION RANK*, including analyst and firm fixed effects. The autocorrelation coefficient is -0.006 with a t-statistic of -0.88.

To implement this, we perform tests based on the state variable being the timing of information arrival. We find that our results are robust to excluding forecasts that are made in the early morning, pre-noon, and post-noon.

We also repeat our main analysis omitting all forecasts in which firms announce earnings in the preceding day (possibly outside of trading hours). If analysts preferentially issue forecasts earlier in the day for firms that announced earnings on the preceding day, then the results might be driven by the increase in accuracy deriving from use of the new public signal that was not yet embedded in the consensus. In Table 5, we find that the results are basically unaffected by omitting the forecasts made following the day of an earnings announcement of a firm.¹⁴ The inferences are identical, and the magnitudes of the coefficients are very similar.

Since all analysts can observe public information when it arrives, the timing of the arrival of public information does not explain the relative forecast accuracy across analysts. Private information can arrive at different times of the day across analysts. However, Reg FD prohibits managers from providing private information to analysts. Given this, we view it as especially unlikely for an analyst to be able consistently to generate private information (for example, by developing relationships with multiple investor relations officers) about different firms *and* for that information to arrive systematically at the beginning of the day for all those firms.

Since we do not have an experiment with a random assignment of forecast ordering for the analysts, we cannot fully rule out all alternative explanations. However, alternative explanations are relatively implausible. A viable alternative explanation must provide both a ranking rule that is neither firm- nor analyst-specific and explain the decline in decision quality relative to self and to peers as the rank increases (later forecast).

In addition to our main results, we also tested the effect of decision fatigue among analysts during earning announcements season. During earning season, analysts are overloaded with

---

¹⁴ The results are also similar when removing forecasts that are made within three and five days after the earnings announcement.

information, so it might seem plausible that decision fatigue effects might be stronger. The higher workload is likely to reduce the amount of time and attention that an analyst devotes to each forecast. However, it is less obvious whether it predicts a stronger decline in accuracy throughout the day.

We repeat all the tests for days following earning announcements (intervals from the day of the forecast to 5 days). In all specifications, we obtain results similar to those for the non-post announcements day sample. A statistical test indicates no significant differences between the two samples.

We also tested if there are cross-sectional differences in the impact of decision fatigue across analysts or across forecasts based on factors that might influence the difficulty of forecasting. We find no cross-sectional difference in decision fatigue based upon the brokerage house size, the experience of the analysts, the existence of managers' guidance, the quality of managers' guidance, or the analysts' company's specific experience. Again, it is plausible that these factors would predict overall level of forecast accuracy or the probability of using heuristics, but not necessarily a greater rate of decline in performance throughout the day.

5. Conclusion

We investigate whether decision fatigue is systematically associated with the forecasting behavior of sell-side security analysts' annual EPS forecasts. Our results suggest that analysts become decision-fatigued during the day, which is consistent with views of cognitive processing developed by Baumeister et al. (1998) and Kahneman and Egan (2011). When mental resources are high, analysts use System 2 thinking and make well-reasoned decisions. However, when mental resources are low, analysts begin to use System 1 thinking and make more intuitive, heuristic decisions. Our archival data test design helps address the potential irreproducibility problems that plague laboratory experiments that are most commonly employed to study ego depletion. We provide a distinct form of evidence about the negative

consequences of decision fatigue as predicted in the psychology literature.

Specifically, we use the number of forecasts an analyst has issued earlier in the same day as a proxy for decision fatigue, and we find that analysts become less accurate as they become more decision-fatigued. We also find that analysts become more heuristic in their forecasting strategies as they become more decision-fatigued; they are more likely to herd toward the consensus forecast, to self-herd by reissuing their own previous outstanding forecast, and to issue a forecast that is rounded to end with a 0 or 5. Finally, we test how the market reacts to these forecasts in relation to the extent that decision fatigue affected the development of the forecast. We find that the stock market’s reaction to a forecast revision is weaker when the issuing analyst is more decision-fatigued.

We can rule out several alternative explanations. First, by controlling for the time of day, we can be confident that we are examining decision fatigue rather than physical fatigue. Second, our difference-in-difference design mitigates firm characteristic effects and focuses instead on the variation in the degree of decision fatigue across analysts in our tests. Third, by removing forecasts that follow earnings announcements, we mitigate the concern that our results are driven by new information.

Some behavioral economics research focuses on cognitive constraints or illusions that are implicitly assumed to be constant over time for any given individual, without regard to exhaustion of cognitive resources. In contrast, our tests focus on how the judgment of economic decision-makers varies as a function of past actions—the decisions made earlier in the day that induce decision fatigue.

Analysts and their firms can potentially benefit from being aware of the influence of decision fatigue on their forecasting behavior and take active measures to minimize its effect. First, firms can implement routines into their workday to minimize the number of less important decisions they make (for example, what to eat and wear) and preserve mental resources for forecast-related decisions. Second, analysts can be assigned to work on the more important forecasts at the beginning of the day, and the firm can progressively increase the

quality monitoring as the number of same-day forecasts increases. Third, small interventions that increase the motivation level (e.g., small incentives, per Baumeister and Vohs 2007), or improve the mood in general (e.g., per Tice et al. 2007) can decrease the impact of decision fatigue.

Our study motivates several directions for future research. First, field evidence on decision fatigue, especially among professionals, is scarce. Future research should explore how decision fatigue influences other agents in the financial system (e.g., retail investors, professional money managers, financial advisors, rating agencies, and market makers) and outside the financial system, especially in professions where mistakes can be very costly (e.g., surgeons, air traffic controllers).

Our study motivates several directions for future research. First, field evidence on decision fatigue, especially among professionals, is scarce. Future research should explore how decision fatigue influences other agents in the financial system (e.g., retail investors, professional money managers, financial advisors, rating agencies, and market makers) and outside the financial system, especially in professions where mistakes can be very costly (e.g., surgeons, air traffic controllers).

Second, future work should explore the sources of decision fatigue, such as personality traits and professional ability. Do incentives or implementation of routines help to mitigate the impact of decision fatigue? The ego depletion hypothesis predicts that incentives would not improve the decision quality of a fatigued individual since the mental resource was exhausted, while the ego preservation hypothesis predicts that incentives would allow the individual to gain access to a mental resource reserve (see Evans et al. 2016 and section 2). This line of research can potentially improve work productivity, influence labor compensation structure.

Third, we find that the market reaction is weaker for high-ranked forecast revision. Future research can explore whether investors' response is influenced by the analyst's decision rank itself or by the use of heuristics such as rounding and herding.

Fourth, to the best of our knowledge, our study is the first to use reissued forecasts. A complete investigation of reissuance activity is in order, including the strategic use of reissuance by analysts, and the effect of reissuance on stock price, volatility, and volume.

References

- Altınkılıç, O., Balashov, V. S., Hansen, R. S., 2013. Are analysts' forecasts informative to the general public? *Management Science* 59, 2550–2565.
- Augenblick, N., Nicholson, S., 2015. Ballot position, choice fatigue, and voter behaviour. *The Review of Economic Studies* 83, 460–480.
- Banerjee, A. V., 1992. A simple model of herd behavior. *The quarterly journal of economics* 107, 797–817.
- Baumeister, R. E., Bratslavsky, E., Muraven, M., Tice, D. M., 1998. Ego depletion: Is the active self a limited resource? *Journal of Personality and Social Psychology* 74, 1252–1265.
- Baumeister, R. F., Tierney, J., 2012. *Willpower: Rediscovering the greatest human strength*. Penguin.
- Baumeister, R. F., Vohs, K. D., 2007. Self-regulation, ego depletion, and motivation. *Social and personality psychology compass* 1, 115–128.
- Baumeister, R. F., Vohs, K. D., 2016. Misguided effort with elusive implications. *Perspectives on Psychological Science* 11, 574–575.
- Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy* 100, 992–1026.
- Bonner, S. E., Walther, B. R., Young, S. M., 2003. Sophistication-related differences in investors' models of the relative accuracy of analysts' forecast revisions. *The Accounting Review* 78, 679–706.
- Bradley, D., Clarke, J., Lee, S., Ornathanalai, C., 2014. Are analysts' recommendations informative? intraday evidence on the impact of time stamp delays. *The Journal of Finance* 69, 645–673.

- Bradshaw, M. T., 2011. Analysts' forecasts: what do we know after decades of work? Working paper, Boston College. Available at <https://ssrn.com/abstract=1880339>.
- Brown, P., Foster, G., Noreen, E., 1985. Security analyst multi-year earnings forecasts and the capital market, vol. 21 of *Studies in accounting research*. American Accounting Association.
- Christophe, S. E., Ferri, M. G., Hsieh, J., 2010. Informed trading before analyst downgrades: Evidence from short sellers. *Journal of Financial Economics* 95, 85–106.
- Clement, M. B., 1999. Analyst forecast accuracy: Do ability, resources, and portfolio complexity matter? *Journal of Accounting and Economics* 27, 285–303.
- Clement, M. B., Tse, S. Y., 2005. Financial analyst characteristics and herding behavior in forecasting. *The Journal of finance* 60, 307–341.
- Cowen, A., Groyberg, B., Healy, P., 2006. Which types of analyst firms are more optimistic? *Journal of Accounting and Economics* 41, 119–146.
- Danziger, S., Levav, J., Avnaim-Pesso, L., 2011. Extraneous factors in judicial decisions. *Proceedings of the National Academy of Sciences* 108, 6889–6892.
- Dechow, P. M., You, H., 2012. Analysts' motives for rounding eps forecasts. *The Accounting Review* 87, 1939–1966.
- Evans, D. R., Boggero, I. A., Segerstrom, S. C., 2016. The nature of self-regulatory fatigue and ego depletion lessons from physical fatigue. *Personality and Social Psychology Review* 20, 291–310.
- Givoly, D., Lakonishok, J., 1979. The information content of financial analysts' forecasts of earnings: Some evidence on semi-strong inefficiency. *Journal of Accounting and Economics* 1, 165–185.

- Gleason, C. A., Lee, C. M., 2003. Analyst forecast revisions and market price discovery. *The Accounting Review* 78, 193–225.
- Gonedes, N. J., Dopuch, N., Penman, S. H., 1976. Disclosure rules, information-production, and capital market equilibrium: The case of forecast disclosure rules. *Journal of Accounting Research* pp. 89–137.
- Groysberg, B., Healy, P. M., 2013. *Wall Street research: Past, present, and future*. Stanford University Press.
- Hagger, M. S., Chatzisarantis, N. L., 2016. Commentary: Misguided effort with elusive implications, and sifting signal from noise with replication science. *Frontiers in psychology* 7, 621.
- Hagger, M. S., Chatzisarantis, N. L., Alberts, H., Anggono, C. O., Batailler, C., Birt, A. R., Brand, R., Brandt, M. J., Brewer, G., Bruyneel, S., et al., 2016. A multilab preregistered replication of the ego-depletion effect. *Perspectives on Psychological Science* 11, 546–573.
- Herrmann, D., Thomas, W. B., 2005. Rounding of analyst forecasts. *The Accounting Review* 80, 805–823.
- Hirshleifer, D., Lim, S. S., Teoh, S. H., 2009. Driven to distraction: Extraneous events and underreaction to earnings news. *The Journal of Finance* 64, 2289–2325.
- Hodge, F. D., 2003. Investors’ perceptions of earnings quality, auditor independence, and the usefulness of audited financial information. *Accounting Horizons* 17, 37–48.
- Hoechle, D., Schaub, N., Schmid, M., 2015. Time stamp errors and the stock price reaction to analyst recommendation and forecast revisions .
- Hong, H., Kubik, J. D., Solomon, A., 2000. Security analysts’ career concerns and herding of earnings forecasts. *The Rand journal of economics* pp. 121–144.

- Jacob, J., Lys, T. Z., Neale, M. A., 1999. Expertise in forecasting performance of security analysts. *Journal of Accounting and Economics* 28, 51–82.
- Kahneman, D., Egan, P., 2011. *Thinking, fast and slow*, vol. 1. Farrar, Straus and Giroux New York.
- Kirk, M., 2011. Research for sale: Determinants and consequences of paid-for analyst research. *Journal of Financial Economics* 100, 182–200.
- Kothari, S., So, E., Verdi, R., 2016. Analysts forecasts and asset pricing: A survey. *Annual Review of Financial Economics* 8, 197–219.
- Kumar, A., 2010. Self-selection and the forecasting abilities of female equity analysts. *Journal of Accounting Research* 48, 393–435.
- Levav, J., Heitmann, M., Herrmann, A., Iyengar, S. S., 2010. Order in product customization decisions: Evidence from field experiments. *Journal of Political Economy* 118, 274–299.
- Lewis, M., 2012. Obama’s way. *The Hive*. Available at <https://www.vanityfair.com/news/2012/10/michael-lewis-profile-barack-obama>.
- 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.
- Michaely, R., Womack, K. L., 1999. Conflict of interest and the credibility of underwriter analyst recommendations. *The Review of Financial Studies* 12, 653–686.
- Mola, S., Guidolin, M., 2009. Affiliated mutual funds and analyst optimism. *Journal of Financial Economics* 93, 108–137.
- Muraven, M., Baumeister, R. F., 2000. Self-regulation and depletion of limited resources: Does self-control resemble a muscle? *Psychological bulletin* 126, 247.

- O'Brien, P. C., Bhushan, R., 1990. Analyst following and institutional ownership. *Journal of Accounting Research* pp. 55–76.
- Ramnath, S., Rock, S., Shane, P., 2008. The financial analyst forecasting literature: A taxonomy with suggestions for further research. *International Journal of Forecasting* 24, 34–75.
- Sripada, C., Kessler, D., Jonides, J., 2016. Sifting signal from noise with replication science. *Perspectives on Psychological Science* 11, 576–578.
- Stanovich, K. E., West, R. F., 2000. Individual differences in reasoning: Implications for the rationality debate? *Behavioral and brain sciences* 23, 645–665.
- Stickel, S. E., 1991. Common stock returns surrounding earnings forecast revisions: More puzzling evidence. *Accounting Review* pp. 402–416.
- Tice, D. M., Baumeister, R. F., Shmueli, D., Muraven, M., 2007. Restoring the self: Positive affect helps improve self-regulation following ego depletion. *Journal of experimental social psychology* 43, 379–384.
- Vohs, K. D., Baumeister, R. F., Schmeichel, B. J., Twenge, J. M., Nelson, N. M., Tice, D. M., 2008. Making choices impairs subsequent self-control: A limited-resource account of decision making, self-regulation, and active initiative. *Journal of Personality and Social Psychology* 94, 883–898.
- Wadhwa, T., 2016. A hedge fund wrote a letter to investors explaining why they should read a classic book about cognitive biases.
- Welch, I., 2000. Herding among security analysts. *Journal of Financial economics* 58, 369–396.
- Zhang, X., 1998. Earnings forecasting and equity valuation under conservative accounting. Doctoral dissertation, Columbia University .

Table 1: Sample Size

Number of Forecasts per day ^a	Number of Analyst-Days ^b	Number of Forecasts ^c
1	255,613	255,613
2	27,975	55,950
3	6,536	19,608
4	2,796	11,184
5	1,559	7,795
6	1,020	6,120
7	766	5,362
8	534	4,272
9	405	3,645
≥ 10	1,326	17,375
Average: 1.3	Total: 298,530	Total: 386,924

^a Number of annual EPS forecasts the analyst has made during day t.

^b Number of distinct analyst-days in which an analyst has made at least one forecast.

^c Number of distinct analyst-day-forecast in the sample.

Table 2: Descriptive Statistics by Decision Order

	Forecast 1	Forecast 2	Forecast 3	Forecast 4	Forecast ≥ 5
<u>Mean</u>					
<i>RELATIVE ACCURACY</i>	0.481	0.392	0.244	0.190	0.107
<i>HERDING</i>	27.8%	30.1%	32.2%	35.3%	37.7%
<i>REISSUE</i>	0.57	0.64	0.74	0.82	0.90
<i>ROUNDING</i>	29.5%	29.8%	29.7%	31.0%	28.3%
<i>TIME OF DAY</i>	3.41	4.32	4.58	4.59	4.58
<i>FIRM EXPERIENCE</i>	0.36	0.36	0.34	0.33	0.36
<i>BROKER SIZE</i>	0.27	0.25	0.23	0.22	0.18
<i>EFFORT</i>	0.58	0.57	0.56	0.56	0.59
<i>FIRMS FOLLOWED</i>	0.42	0.45	0.45	0.44	0.48
<i>FORECAST AGE</i>	0.50	0.51	0.50	0.49	0.49
<i>NUMEST</i>	2.35	2.44	2.55	2.62	2.69
<i>MVE</i>	7.95	8.03	8.13	8.26	8.19
<i>ROA</i>	2.7%	3.4%	4.3%	4.5%	4.5%
<i>SALES GROWTH</i>	0.3%	0.2%	0.2%	0.2%	0.2%
<i>R&D</i>	62.6%	60.5%	57.0%	52.3%	37.7%
<i>INTANGIBLE ASSETS</i>	17.6%	15.4%	13.2%	11.5%	7.8%
<i>ADVERTISING</i>	3.76	3.90	4.10	4.33	4.26
<i>CFF</i>	0.19	0.13	0.09	0.08	0.13
<i>STD RET</i>	0.15	0.15	0.15	0.16	0.15
<i>E/P</i>	-0.019	-0.010	0.003	0.000	0.008
<i>VOLUME</i>	0.331	0.332	0.331	0.334	0.327
<u>Median</u>					
<i>RELATIVE ACCURACY</i>	0.314	0.274	0.200	0.152	0.139
<i>HERDING</i>	0.00%	0.00%	0.00%	0.00%	0.00%
<i>REISSUE</i>	100.00%	100.00%	100.00%	100.00%	100.00%
<i>ROUNDING</i>	0.00%	0.00%	0.00%	0.00%	0.00%
<i>TIME OF DAY</i>	3.00	5.00	5.00	5.00	5.00
<i>FIRM EXPERIENCE</i>	0.250	0.263	0.250	0.238	0.273
<i>BROKER SIZE</i>	0.149	0.131	0.111	0.098	0.050
<i>EFFORT</i>	0.571	0.571	0.556	0.556	0.600
<i>FIRMS FOLLOWED</i>	0.375	0.412	0.405	0.391	0.440
<i>FORECAST AGE</i>	0.513	0.515	0.514	0.507	0.491
<i>NUMEST</i>	2.4	2.49	2.64	2.74	2.77
<i>MVE</i>	7.87	7.96	8.05	8.24	8.14
<i>ROA</i>	5.13%	5.35%	5.63%	5.78%	5.91%
<i>SALES GROWTH</i>	0.03%	0.03%	0.02%	0.02%	0.03%
<i>R&D</i>	100.00%	100.00%	100.00%	100.00%	0.00%
<i>INTANGIBLE ASSETS</i>	10.80%	7.65%	5.11%	4.04%	1.41%
<i>ADVERTISING</i>	3.78	3.94	4.15	4.26	4.25
<i>CFF</i>	-0.148	-0.153	-0.154	-0.146	-0.061
<i>STD RET</i>	0.141	0.141	0.141	0.142	0.141
<i>E/P</i>	0.043	0.046	0.049	0.051	0.054
<i>VOLUME</i>	0.328	0.328	0.328	0.328	0.319

The table presents mean and median of our variable of interest by number of the forecast made by the analyst for the given day. The sample includes all annual EPS forecasts on days when the analyst only issued forecasts between the working hours of 9:00 a.m. and 7:00 p.m between the years 2002-2015. Variable definitions are in Appendix A.

Table 3: Relative Accuracy and Decision Fatigue

	(1)	(2)	(3)	(4)	(5)	(6)
<i>DECISION RANK</i>	-0.303*** (-22.90)	-0.225*** (-16.71)	-0.181*** (-9.60)	-0.169*** (-9.11)	-0.042** (-2.08)	-0.067*** (-2.85)
<i>TIME OF DAY</i>		-0.007*** (-6.12)		-0.006*** (-4.65)		0.006 (1.17)
<i>FIRM EXPERIENCE</i>		0.137*** (14.08)		0.046*** (3.17)		0.041* (1.74)
<i>BROKER SIZE</i>		0.038*** (3.26)		0.033 (1.18)		0.006 (0.11)
<i>EFFORT</i>		0.024** (2.22)		-0.088*** (-6.00)		-0.101*** (-3.71)
<i>FIRMS FOLLOWED</i>		0.004 (0.32)		0.016 (0.80)		0.022 (0.54)
<i>FORECAST AGE</i>		-0.183*** (-15.93)		-0.170*** (-12.40)		0.044 (1.16)
<i>NUMEST</i>		-0.232*** (-46.10)		-0.184*** (-25.51)		-0.088*** (-6.65)
Constant	0.693*** (64.10)	1.224*** (61.68)	0.599*** (41.20)	1.151*** (39.23)	0.491*** (31.76)	0.706*** (12.85)
Fixed Effects	N	N	Analyst	Analyst	Analyst-day	Analyst-day
Adjusted R-squared	0.001	0.010	0.045	0.049	0.398	0.398
Observations	386,924	386,924	386,924	386,924	386,924	386,924

The dependent variable is as follows: $RELATIVE\ ACCURACY_{i,j,t}$ is analyst i 's EPS forecast error of company j at day t . This EPS forecast error is compared to the median EPS forecast error for all analysts issuing EPS forecast error for company j up until day t (consensus). The relative accuracy is standardized across firms by deflating the standard deviation of EPS forecasts error across all analysts who cover the firm. The independent variables are as follows: *DECISION RANK* is the log value of the number of forecasts an analyst has made before the forecast being evaluated, plus 1. Definitions of the control variables are provided in Appendix A. t-statistics are provided in parentheses with heteroskedastic-consistent standard errors clustered at the analyst level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 4: Herding and Decision Fatigue

	(1)	(2)	(3)	(4)	(5)	(6)
<i>DECISION RANK</i>	0.348*** (13.03)	0.267*** (10.05)	0.167*** (7.49)	0.162*** (7.18)	0.082** (2.40)	0.086** (2.04)
<i>TIME OF DAY</i>		0.008*** (3.87)		0.003* (1.73)		0.003 (0.33)
<i>FIRM EXPERIENCE</i>		-0.008 (-0.37)		0.170*** (5.39)		0.009 (0.25)
<i>BROKER SIZE</i>		-0.058** (-2.18)		0.093** (2.18)		0.216** (2.48)
<i>EFFORT</i>		0.006 (0.23)		0.141*** (6.09)		0.074 (1.60)
<i>FIRMS FOLLOWED</i>		0.064* (1.83)		-0.006 (-0.19)		0.060 (0.93)
<i>FORECAST AGE</i>		-0.170*** (-8.76)		-0.137*** (-7.96)		-0.243*** (-3.45)
<i>NUMEST</i>		0.217*** (21.30)		0.182*** (14.13)		0.145*** (6.89)
Constant	-1.199*** (-53.41)	-1.620*** (-41.47)				
Fixed Effects	N	N	Analyst-Firm	Analyst-Firm	Analyst-day	Analyst-day
Pseudo R-squared	0.0009	0.0048	0.0002	0.0016	0.0001	0.0013
Observations	324,456	324,456	263,839	263,839	61,276	61,276

The dependent variable, $HERDING_{i,j,t}$, is a binary variable with a value of 1 if analyst i forecast of company j at time t is between the consensus forecast at time t and his own previous forecast, and 0 otherwise. The independent variables are as follows: *DECISION RANK* is the log value of the number of forecasts an analyst has made before the forecast being evaluated, plus 1. Definitions of the control variables are provided in Appendix A. z-statistics are provided in parentheses with heteroskedastic-consistent standard errors clustered at the analyst level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 5: Reissuance of a Previous Outstanding Forecast and Decision Fatigue

	(1)	(2)	(3)	(4)	(5)	(6)
<i>DECISION RANK</i>	1.230*** (28.98)	1.151*** (27.90)	1.419*** (117.95)	1.349*** (110.68)	1.845*** (57.79)	1.927*** (39.00)
<i>TIME OF DAY</i>		0.022*** (6.20)		0.027*** (26.48)		-0.014** (-2.30)
<i>FIRM EXPERIENCE</i>		0.089*** (3.92)		0.230*** (12.57)		0.052 (1.35)
<i>BROKER SIZE</i>		0.558*** (19.74)		0.434*** (17.81)		-0.048 (-0.59)
<i>EFFORT</i>		0.088*** (3.30)		0.105*** (8.19)		-0.017 (-0.40)
<i>FIRMS FOLLOWED</i>		-0.080*** (-2.99)		-0.045** (-2.48)		-0.053 (-0.85)
<i>FORECAST AGE</i>		-0.798*** (-61.24)		-0.893*** (-84.97)		-0.972*** (-17.75)
<i>NUMEST</i>		0.106*** (8.39)		0.183*** (21.87)		0.113*** (5.60)
Constant	-0.571*** (-17.30)	-0.673*** (-11.16)				
Fixed Effects	N	N	Analyst-Firm	Analyst-Firm	Analyst-day	Analyst-day
Pseudo R-squared	0.02	0.03	0.02	0.04	0.10	0.11
Observations	696,884	696,884	653,156	653,156	52,252	52,252

The dependent variable, $REISSUE_{i,j,t}$, is a binary variable with a value of 1 if analyst i forecast of company j at time t is the reissuance of her own previous forecast, and 0 otherwise. The independent variables are as follows: *DECISION RANK* is the log value of the number of forecasts an analyst has made before the forecast being evaluated, plus 1. Definitions of the control variables are provided in Appendix A. z-statistics are provided in parentheses with heteroskedastic-consistent standard errors clustered at the analyst level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 6: Rounding and Decision Fatigue

	(1)	(2)	(3)	(4)	(5)	(6)
<i>DECISION RANK</i>	-0.046 (-0.68)	0.016 (0.24)	0.037 (1.27)	0.052* (1.75)	0.061 (1.31)	0.136** (2.36)
<i>TIME OF DAY</i>		0.001 (0.24)		-0.005** (-1.97)		-0.027** (-2.38)
<i>HERDING</i>		-0.012 (-0.91)		0.002 (0.14)		0.049* (1.79)
<i>FIRM EXPERIENCE</i>		0.260*** (5.16)		0.045 (0.99)		0.011 (0.20)
<i>BROKER SIZE</i>		0.066 (1.15)		0.067 (1.16)		-0.285** (-2.05)
<i>EFFORT</i>		-0.531*** (-9.90)		-0.137*** (-4.48)		-0.066 (-1.00)
<i>FIRMS FOLLOWED</i>		-0.007 (-0.12)		-0.000 (-0.01)		0.073 (0.78)
<i>FORECAST AGE</i>		0.628*** (22.15)		0.778*** (34.33)		0.773*** (7.75)
<i>NUMEST</i>		-0.146*** (-7.16)		-0.010 (-0.62)		-0.013 (-0.41)
Constant	-0.742*** (-15.07)	-0.518*** (-5.41)				
Fixed Effects	N	N	Analyst-Firm	Analyst-Firm	Analyst-Day	Analyst-Day
Pseudo R-squared	0.0000	0.0097	0.0000	0.0086	0.0001	0.0029
Observations	205,228	205,228	165,802	165,802	34,747	34,747

The dependent variable, $Rounding_{i,j,t}$, is a binary variable with a value of 1 if analyst i 's forecast of company j at time t ends with a 0 or 5 in the penny digit, and 0 otherwise. The independent variables are as follows: *DECISION RANK* is the log value of the number of forecasts an analyst has made before the forecast being evaluated, plus 1. Definitions of the control variables are provided in Appendix A. z-statistics are provided in parentheses with heteroskedastic-consistent standard errors clustered at the analyst level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 7: Stock Market Reaction to Analyst Forecast Revision and Decision Fatigue

	(1)	(2)	(3)	(4)	(5)	(6)
<i>DECISION RANK</i>	0.002*** (2.68)	0.001** (2.24)	0.001 (1.06)	0.001 (0.94)	-0.000 (-0.01)	-0.001 (-1.40)
<i>FORECAST REVISION</i>	0.017*** (48.77)	0.014*** (24.28)	0.017*** (44.07)	0.014*** (21.67)	0.011*** (21.00)	0.010*** (10.05)
<i>DECISION RANK*</i>	-0.007*** (-15.02)	-0.006*** (-12.57)	-0.007*** (-13.20)	-0.005*** (-11.02)	-0.002*** (-4.63)	-0.001** (-2.27)
<i>FORECAST REVISION</i>						
Controls	N	Y	N	Y	N	Y
Controls*						
FORECAST REVISION	N	Y	N	Y	N	Y
Fixed Effects	N	N	Analyst-Firm	Analyst-Firm	Analyst-day	Analyst-day
Adjusted R-squared	0.117	0.122	0.168	0.172	0.565	0.568
Observations	324,456	324,456	324,456	324,456	324,456	324,456

The dependent variable $CAR_{i,j,t}$ is the 3-day market-adjusted excess return for firm j centered on the forecast revision issued by analyst i at time t . The independent variables are as follows: *DECISION RANK* is the log value of the number of forecasts an analyst has made before the forecast being evaluated, plus 1. *FORECAST REVISION* is a measure of the difference between the current annual earnings forecast for analyst i who follows firm j in time t and the annual earnings forecast issued immediately before current annual earnings forecast, scaled by the standard deviation of forecasts of all analysts who cover firm j in time t . Definitions of the control variables are provided in Appendix A. t-statistics are provided in parentheses with heteroskedastic-consistent standard errors clustered at the analyst level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 8: Decision Fatigue with Analyst-Firm-Year Fixed Effects

Dependent:	(1) Accuracy	(2) Herding	(3) Reissue	(4) Rounding
<i>DECISION RANK</i>	-0.127*** (-7.55)	0.113*** (4.03)	1.328*** (96.48)	0.017 (0.46)
<i>TIME OF DAY</i>	-0.000 (-0.13)	0.001 (0.29)	0.044*** (38.55)	-0.004 (-1.27)
<i>BROKER SIZE</i>	-0.042 (-0.28)	0.861*** (4.12)	-0.344*** (-3.07)	0.284 (0.87)
<i>FORECAST AGE</i>	-0.303*** (-20.40)	-0.099*** (-4.65)	-1.190*** (-90.65)	0.708*** (25.25)
<i>NUMEST</i>	-0.276*** (-20.86)	0.257*** (13.76)	0.288*** (24.96)	-0.020 (-0.86)
<i>HERDING</i>				0.011 (0.67)
Constant	1.376*** (26.16)			
Fixed Effects	Analyst-Firm-Year	Analyst-Firm-Year	Analyst-Firm-Year	Analyst-Firm-Year
Adj/Psdo R-sqr	0.212	0.002	0.053	0.009
Observations	386,924	157,899	525,419	97,983

The dependent variables Relative Accuracy, Herding, Reissue, and Rounding are described in Tables 3, 4, 5, and 6 respectively. Definitions of the control variables are provided in Appendix A. t-statistics (z-statistics) are provided in parentheses with heteroskedastic-consistent standard errors clustered at the analyst level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 9: Forecasting Behavior and Decision Fatigue, Excluding Forecasts Following an Earnings Announcement

Dependent:	(1) Accuracy	(2) Herding	(3) Reissue	(4) Rounding
<i>DECISION RANK</i>	-0.058** (-2.42)	0.111** (2.42)	2.298*** (39.72)	0.105** (2.00)
<i>TIME OF DAY</i>	-0.001 (-0.13)	-0.001 (-0.11)	-0.076*** (-10.36)	-0.022* (-1.80)
<i>FIRM EXPERIENCE</i>	0.051** (2.13)	0.042 (1.02)	0.103** (2.23)	0.001 (0.03)
<i>BROKER SIZE</i>	0.055 (0.95)	0.275*** (2.84)	0.116 (1.16)	-0.308** (-2.31)
<i>EFFORT</i>	-0.076*** (-2.72)	0.013 (0.27)	-0.085 (-1.64)	-0.084 (-1.39)
<i>FIRMS FOLLOWED</i>	0.029 (0.70)	0.033 (0.47)	-0.096 (-1.28)	-0.008 (-0.09)
<i>FORECAST AGE</i>	0.084** (2.13)	-0.193** (-2.47)	-0.987*** (-15.83)	0.693*** (8.92)
<i>NUMEST</i>	-0.050*** (-3.66)	0.137*** (5.79)	0.062** (2.55)	-0.033 (-1.14)
Constant	0.497*** (8.52)			
Fixed Effects	Analyst-day	Analyst-day	Analyst-day	Analyst-day
Adj/Psdo R-sqr	0.441	0.00110	0.115	0.00307
Observations	313,841	53,393	37,707	40,268

The sample used in this Table does not include forecasts that are made following the day after an earnings announcement of a firm. The dependent variables are as follows: $RELATIVE\ ACCURACY_{i,j,t}$ is analyst i 's EPS forecast error of company j at day t . This EPS forecast error is compared to the median EPS forecast error for all analysts issuing EPS forecast error for company j up until day t (consensus). The relative accuracy is standardized across firms by deflating the standard deviation of EPS forecasts error across all analysts who cover the firm. $HERDING_{i,j,t}$ is a binary variable with a value of 1 if analyst i 's forecast of company j at time t is between the consensus forecast at time t and her own previous forecast, and 0 otherwise. $REISSUE_{i,j,t}$ is a binary variable with a value of 1 if analyst i 's forecast of company j at time t is the reissuance of her own previous forecast, and 0 otherwise. The independent variables are as follows: *DECISION RANK* is the log value of the number of forecasts an analyst has made before the forecast being evaluated, plus 1. Definitions of the control variables are provided in Appendix A. t-statistics (z-statistics) are provided in parentheses with heteroskedastic-consistent standard errors clustered at the analyst level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Appendix A. Variable Definitions

Variable Name	Description
<i>ADVERTISING</i>	Log value of advertising expense plus 1.
<i>BROKER SIZE</i>	A measure of the size of analyst i 's brokerage house. It is calculated as the number of analysts employed by the brokerage that employs analyst i following firm j in year t minus the minimum number of analysts employed by brokerages for analysts who follow firm j in year t , with this difference scaled by the range of brokerage house sizes for analysts who follow firm j in year t .
<i>CAR</i>	The 3-day market-adjusted excess return for firm j centered on the forecast revision issued by analyst i at time t .
<i>CFF</i>	Cash flows from financing at year t deflated by the absolute value of operating cash flows at year t .
<i>DECISION RANK</i>	The log value of the number of forecasts an analyst has made before the forecast being evaluated, plus 1.
<i>E/P</i>	Earnings per share (basic) at the end of year t deflated by price per share at the end of year t .
<i>EFFORT</i>	A measure of analyst i 's effort in forecasting firm j . It is calculated as the number of forecasts issued by analyst i following firm j in year t minus the minimum number of forecasts issued by analysts who follow firm j in year t , with this difference scaled by the range of forecasts issued by analysts who follow firm j in year t .

FIRM EXPERIENCE

A measure of analyst i 's firm-specific experience. It is calculated as the number of years of firm-specific experience for analyst i following firm j in year t minus the minimum number of years of firm-specific experience for analysts who follow firm j in year t , with this difference scaled by the range of years of firm-specific experience for analysts who follow firm j in year t .

FIRMS FOLLOWED

A measure of the number of companies that analyst i follows in year t . It is calculated as the number of companies followed by analyst i following firm j in year t minus the minimum number of companies followed by analysts who follow firm j in year t , with this difference scaled by the range in the number of companies followed by the analysts who follow firm j in year t .

FORECAST AGE

A measure of the time from the forecast date to the earnings announcement. It is calculated as the number of days from the forecast date to the date of the earnings announcement for analyst i in year t minus the minimum number of days from the forecast date to the date of the earnings announcement for analysts who follow firm j in year t , with this difference scaled by the range of days from the forecast date to the date of the earnings announcement for analysts who follow firm j in year t .

FORECAST REVISION

A measure of the difference between the current annual earnings forecast for analyst i following firm j in time t and the annual earnings forecast issued immediately before the current annual earnings forecast, scaled by the standard deviation of forecasts of all analysts who cover firm j in time t .

<i>HERDING</i>	A dummy variable that receives the value of 1 for forecasts that are between the analyst's own prior forecast and the consensus forecast, and 0 otherwise.
<i>INTANGIBLE ASSETS</i>	Intangible assets at the end of the fiscal year deflated by total assets at the end of the fiscal year.
<i>MVE</i>	Price per share at the end of the fiscal year multiplied by number of shares at the end of the fiscal year.
<i>NUMEST</i>	The number of analysts who cover firm j at time t .
<i>R&D</i>	A dummy variable that receives the value of one if the firm has an R&D expense at the end of the year and zero otherwise.
<i>REISSUE</i>	A dummy variable that takes the value of one if a forecast is reissued (self-herding) and zero otherwise.
<i>RELATIVE ACCURACY</i>	A measure of analyst i 's EPS forecast error for company j at time t subtracted from the median EPS forecast error for all analysts who cover firm j within the same 90 days. This difference is standardized across firms by dividing it by the standard deviation of EPS forecast errors across all analysts who cover firm j at time t .
<i>ROA</i>	Income before extraordinary items deflated by total assets at the end of the fiscal year.
<i>SALES GROWTH</i>	The value of (sales in fiscal year t minus sales in fiscal year $t - 1$), deflated by sales in fiscal year $t - 1$.
<i>STD RET</i>	Standard deviation of monthly returns in year t .

TIME OF DAY

An ordinal measure that receives the value of 1 for the first hour of the workday (9:00 a.m.), the value of 2 for the second hour of the workday (10:00 a.m.), and so on.

VOLUME

The mean monthly share volume for year t deflated by market value of equity.