

Security Analysts' Career Concerns and Herding of Earnings Forecasts

Harrison Hong
Stanford Graduate School of Business

Jeffrey D. Kubik
Syracuse University

Amit Solomon
MIT Department of Economics

First Draft: May 1997

This Draft: July 1998

Abstract

Several theories of reputation and herding (see, e.g., Scharfstein and Stein (1990)) suggest that herding among agents should vary with career concerns. Our goal in this paper is to document whether such a link exists in the labor market for security analysts. Specifically, we look at the relationship between an analyst's job tenure (a proxy for career concerns) and various measures of stock earnings forecast performance. We establish the following key results. (1) Older analysts are more likely to produce earnings forecasts of firms before younger ones. (2) Their forecasts also deviate more from the consensus forecast than their younger counterparts. We argue that these results are consistent with some reputational models of herding. We also establish a number of auxiliary findings regarding the relationship between forecast accuracy and frequency of forecast revisions with job tenure.

We are grateful to Bengt Holmstrom, Paul Joskow and the MIT Department of Economics for financial support to purchase data for this paper. We are very thankful for extensive comments from Glenn Ellison, Leora Friedberg, and Jeff Zwiebel. Helpful suggestions from Josh Angrist, Dwight Crane, Denis Gromb, Bengt Holmstrom, Jan Mart-Smith, Andre Perrold, Matt Siegel, Jeremy Stein, Joanne Tao, Sarah Tasker, and Jiang Wang are also appreciated.

1. Introduction

A number of recent theoretical studies have examined the relationship between reputation and herding (see, e.g., Scharfstein and Stein (1990), Zwiebel (1995), and Prendergast and Stole (1996)).¹ These models regard reputation as arising from learning over time about some exogenous characteristics of agents, such as ability, through their observed behavior. Reputation affects agents' decisions when they adjust their behavior to influence the data others use to learn about their ability. Such concern for reputation can lead agents to copy the actions of others, i.e., to herd.

These studies suggest an interesting link between the propensity of agents to herd and reputation at various stages of their careers. The idea of an agent managing her reputation because it may affect her wages later in her career dates back to Fama (1980) and Lazear and Rosen (1981). Such "career concerns" are formally modelled by Holmstrom (1982). Unlike these earlier models, which tend to deal with the decisions of a single agent, the more recent multi-agent models of reputation can produce a link between career concerns and herding behavior.

Even though there is a strong theoretical interest in the economics of reputation and herding, there has been relatively little empirical work analyzing the potential connection between career concerns and herding. The goal of this paper is to document this relationship in the labor market for security analysts. Producing earnings forecasts of firms is one of the principal responsibilities of security analysts. These earnings estimates are easily observable and comparable across analysts. In this paper, we examine the relationship between various measures of herding in earnings forecasts and analyst job tenure, a standard proxy for career concerns.

¹ There are several other theoretical papers that analyze herding in specific labor markets. Among these are Trueman (1994), Huddart (1996) and Graham (1998), who consider the labor markets for security and investment analysts. These papers extend the themes of herding developed in Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992).

To be clear, our goal is not to distinguish between various models of reputation and herding. In fact, several existing theories imply a non-zero correlation between herding and job tenure. For a given theory, the sign of the correlation depends on a number of factors. For example, consider Scharfstein and Stein (1990), one of the first treatments of reputation induced herding and the model most useful when thinking about analysts' forecast performance. Reputation leads to herding due to what they call a "sharing-the-blame effect." It is more costly to be wrong and to do something different from what others do than to be wrong just like everyone else. On the one hand, as an analyst's career progresses, there is apt to be less uncertainty about her ability, which should then reduce her incentive to herd. On the other hand, a more experienced analyst may forecast closer to the herd since she has more to lose if she is wrong because it is more costly to switch professions when old than when young. Whether analysts herd more or less with job tenure depends on a tradeoff of these two effects.

By quantifying the relationship between herding and job tenure for security analysts, we can better place these reputation theories in some real world context and understand how important these career concern effects are in practice. As we will argue below, the labor market for security analysts shares many of the characteristics of labor markets considered in the models of reputation and herding. So our findings for the security analyst market could be representative of the association between career concerns and herding in similar labor markets. Additionally, our findings can help to inform and refine future theoretical studies of reputation and herding.

The labor market for security analysts is an ideal setting to test for the link between career concerns and herding. First, there are readily available measures of job performance and job tenure that are rarely simultaneously available in other labor markets. Also, reputation seems, at least anecdotally, to be important for security analysts. For instance, investment banks often advertise the earnings forecast performance of their analysts, such as the timeliness of their estimates, and the profiles of many successful forecasters are published in newspapers such as the *Wall Street Journal*

and financial magazines such as *Fortune* and the *Institutional Investor*. This financial press often suggests that there is some kind of forecasting ability and that analysts are judged relative to the performance of their peers (see, e.g., Stickel (1992)). In addition to the above evidence on ability and relative performance, there is also anecdotal evidence of herd-like behavior in the labor market for security analysts (see, e.g., Nocera (1997)).

Before we can test for the link between career concerns and herding, we need to construct measures of herding in earnings forecasts. From the theoretical literature, we take the following to be consistent with less herding. First, an analyst is less likely to be herding if she issues forecasts before other analysts (i.e., issues more timely forecasts). Second, an analyst is less likely to be herding if her forecast deviates more from the consensus forecast (the average earnings estimate of other analysts covering the firm) than other analysts. Hence, our empirical strategy is to examine regressions of these measures of herding in earnings forecasts on job tenure, controlling for a number of other factors that affect forecast performance. The null hypothesis is that there is no relationship between job tenure and these measures of herding.

To preview, our main results are as follows. First, older security analysts are more likely to produce forecasts before younger ones. (We refer to the age of an analyst and her job tenure interchangeably.) If this finding is indeed due to reputation and herding, we would also expect to find that older analysts' forecasts deviate more from the consensus than younger analysts. Otherwise, interpreting the timeliness result would be problematic. Indeed, we find that older analysts do deviate more from the consensus forecast than younger analysts, providing a comforting confirmation of our first finding. We argue below that these two findings in tandem strongly suggest that older analysts herd less than younger analysts. Finally, we also estimate the relationship between two additional forecast performance measures and job tenure: forecast accuracy and frequency of revisions of forecasts. We find some evidence that these two measures also vary with job tenure. These findings

do not fundamentally alter our interpretation of the previous findings but may provide additional empirical restrictions for future theories interested in the forces shaping forecast performance.

It is important to note the limitations of our analysis. We do not claim that some combination of non-reputation induced herding explanations cannot explain these findings. In fact, we will consider a number of such explanations below; however, we will ultimately argue that reputation-based herding explanations represent the more parsimonious description of the aforementioned results. On these grounds, we view our findings as suggestive of reputation-based herding mechanisms at work in the submission of earnings forecasts.

There have been a number of other empirical studies of herding, but few attempts to document a relationship with career concerns. The market for professional forecasts has attracted many herding papers. For example, Welch (1996) finds some evidence for herding in analyst recommendations on whether to buy or sell stocks. However, Ehrbeck and Waldman (1996) conclude that empirical patterns in 3-month T-bill forecasts do not support simple reputation driven herding models but, instead, are consistent with behavioral hypotheses. Other evidence of herding among pension funds, mutual funds, and institutional investors has been documented by Lakonishok, Shleifer, Thaler and Vishny (1991), Peles (1993), Grinblatt, Titman and Wermers (1995), Wermers (1995), Falkenstein (1996), Nofsinger and Sias (1996) and Wylie (1996). These papers, however, do not focus on the relationship between herding and the career concerns of agents.

There are a few that do. The study most closely related to our work is Lamont (1995). His is the first attempt to document a link between career concerns and herding. He considers the forecast performance of a subset of GNP (and other macro-economic variables) forecasters who are featured in various issues of *Business Week*. Consistent with our results, he finds that older analysts herd less. Unlike his sample of forecasters, ours is a relatively comprehensive set of security analysts. Additionally, his paper only examines one measure of herding, deviation from consensus, as opposed to

ours, which also considers the timing of forecasts. Also, Chevalier and Ellison (1997) report that younger mutual fund managers tend to herd more in the sense that they take on lower unsystematic risk and deviate less from the objectives of their mutual funds than their older counterparts. Finally, Graham (1998) finds that there does not appear to be a relationship between herding and the age of analysts who produce investment newsletters.

The remainder of our paper proceeds as follows. In Section 2, we describe our data and analyze in detail our measures of herding. Section 3 explains our empirical methodology. We discuss our main findings in Section 4. In Section 5, we explore our auxiliary findings and conclude in Section 6.

2. Measures of Herding

Our data come from the Institutional Brokers Estimate System (I/B/E/S) database. I/B/E/S gathers the earnings forecasts of companies throughout the world from thousands of individual security analysts. We use the I/B/E/S Detail Earnings Estimate History File, which contains earnings forecasts of U.S. companies between 1983 and 1996.² During this period, the data consist of the earnings estimates of 8,421 analysts covering 4,527 firms. Before constructing our measures of herding, we establish several key facts about our sample of analysts in Table 1. The goal in presenting these descriptive statistics is twofold. First, we want to demonstrate that our sample is a reasonable one with which to measure the link between career concerns and herding. Second, we want to highlight the features of the data that allow us to estimate this relationship.

² I/B/E/S also distributes the Historical Summary File, which contains more years of data (back to 1976) than the Detail Earnings Estimate History File. However, only the latter contains information on the forecast histories of individual analysts.

First, we consider the distribution of the number of years an analyst is in the I/B/E/S sample. On average, we observe an analyst in the data for a little over four years, with a standard deviation of about 3.7 years. Some are in the database only one year (the 10th percentile of the distribution), either because they quickly left the profession or they only started providing earnings forecasts in 1996. However, a number of analysts are in the sample for the entire fourteen-year period. The 90th percentile of the distribution is ten years. Hence, the sample contains a high proportion of analysts whose forecast performance we can track for many years.

Next, we examine the distribution of the number of analysts that cover a firm in a year. Note that the mean number of analysts covering a firm is 18.59 with a standard deviation of 10.57. The fact that many analysts cover a stock means that we can construct a consensus forecast for a firm in a year by taking the average of the forecasts of all analysts following the firm.

In the third row of Table 1, we consider the distribution of the number of stocks an analyst follows in a year. Generally, analysts tend to specialize and cover firms in the same industry. Hence, they tend to cover more than one stock at any point in time. On average, an analyst follows 9.5 firms in a year (with a standard deviation of about 4.7 firms). We also measure the distribution of the number of years an analyst follows an individual firm. Note that the mean number of years an analyst follows a stock is 2.73, which is smaller than the mean number of years an analyst is in the sample. In other words, analysts tend to switch the firms that they cover as they progress in their jobs. To make this point even more strongly, compare the 90th percentile of the distribution of the number of years an analyst is in the sample (10 years) to the 90th percentile of the distribution of the number of years an analyst follows a stock (6 years). We need to consider this switching in the firms that an analyst covers when measuring how an analyst's forecast behavior changes over time.

With these important features of our sample in mind, we proceed to construct various measures of earnings forecast herding.

2.1. Timing of Forecasts

Perhaps the most intuitive measure of herding is the order in which analysts produce earnings forecasts for a stock (or the timing of forecasts). Certainly, analysts who generate forecasts before other analysts are less likely to be herding than those who issue forecasts after other analysts. In this spirit, our first measure of herding is whether an analyst was the first person to issue an earnings estimate for a stock for a given fiscal year. We assume that an analyst who issues the first forecast of the fiscal year is less likely to be herding than those who do not.

More specifically, we construct our first herding measure as follows. For each firm that an analyst follows in a fiscal year, we identify the date of the first forecast for that analyst. We compare that date to the dates of the first forecasts of the other analysts that also follow that stock to determine the analyst that produced the first forecast for the firm's earnings that year.³ Then, we can determine the relationship between an analyst's tenure and the probability that she moves first.

Because the first earnings estimate of a firm can sometimes be made a few years before the fiscal year being considered, we must examine only analysts who had the opportunity to produce the initial estimate for a firm in a given year. Therefore, the sample is restricted to analysts who have at least four years of experience following the firm. If we were to use the entire sample, younger analysts would be very unlikely to have produced the first forecast for a firm, not because of herding behavior, but because they were not following the firm or even in the profession at the time the first forecast was

³ I/B/E/S does not report the date that an analyst produces an earnings forecast; rather, it provides the date that the analyst contacted I/B/E/S about the earnings estimate. We assume that analysts report their forecasts quickly to I/B/E/S as part of their effort to advertise their work. Occasionally, more than one analyst reports the first earnings forecast on the same date. In such circumstances, all analysts that report a forecast that day are considered to be the first mover.

produced. Including such analysts in the analysis would bias our estimates toward finding that more experienced analysts are more likely to lead.⁴

2.2. Deviation from Consensus

Another natural herding measure is to compare an analyst's forecast to the consensus forecast. Our measure of the consensus forecast is just the average of the forecasts of other analysts following the same firm. An implication of most reputation models is that an analyst who deviates more from the consensus is less likely to be herding than someone who deviates less from the consensus.

What makes this comparison not entirely obvious, however, is that analysts often make many revisions of an earnings forecast for a company over time with no pre-set schedule as they obtain new information about the firm. Our procedure is to examine the latest earnings estimate made by an analyst between January 1st and July 1st of the year being considered (generally six months before the fiscal year ends for most firms). For instance, consider the earnings forecasts made by an analyst for a firm's 1996 fiscal year end earnings. We collect the earnings forecast of an analyst following this firm that is closest to but before July 1, 1996, but we do not include the forecasts of analysts who did not produce a forecast within this six month period from January to July. It is from forecasts within this window that we calculate the consensus forecast for the firm in 1996 and the deviation from consensus measures.

The July 1st date is, for the most part, arbitrarily chosen. Later, we will check the sensitivity of our results to other cutoff dates. It is important that we examine forecasts of analysts made within a

⁴ Another potential problem with our sample is that for early years of our I/B/E/S dataset, we cannot be sure that the first forecast of a firm is during the sample period. For example, when determining who produced the initial forecast of a firm's 1984 earnings, we can only go back to the 1983 estimates. If the first estimates were produced before 1983, then we would miss them. However, this concern is not a problem for our analysis since

specific time period and not all forecasts made by analysts before July 1st. For example, an analyst may produce a forecast for a firm's 1996 fiscal year earnings a few years ahead of time. If, shortly after that early forecast, the analyst stops covering the firm, then that distant forecast would be the analyst's latest forecast if we examined all estimates before July 1st of the fiscal year being considered. However, that estimate would be out of date. If the analyst continued following the firm, she would have probably updated her forecast over time. Therefore, we only want to examine forecasts within a rather narrow window.

To make the above discussion more concrete, we introduce some notation to summarize how we construct this herding measure. Let $F_{i,j,t}$ be the most recent (dollar) earnings per share (EPS) estimate issued by analyst i on stock j before July 1st in year t . Let $\bar{F}_{-i,j,t} = \frac{1}{n} \sum_{m \in -i} F_{m,j,t}$, where $-i$ is the set of all analysts other than analyst i that produce an earnings estimate for stock j in year t and n is the number of analysts in $-i$. Hence, $\bar{F}_{-i,j,t}$ is a measure of the consensus forecast, the average of most recent forecasts (before July 1st) made by all other analysts except analyst i following stock j in year t . We only examine stocks with significant analyst coverage. At least *ten* analysts must be covering a stock for it to be included in our analysis.⁵ Then an intuitive measure of herding is just the absolute value of the difference between $F_{i,j,t}$ and $\bar{F}_{-i,j,t}$:

$$\text{Deviation from Consensus}_{i,j,t} = \left| F_{i,j,t} - \bar{F}_{-i,j,t} \right|. \quad (1)$$

we restrict our sample to analysts with at least four years of experience in the I/B/E/S sample. The first fiscal year considered by these analysts cannot be before 1987.

⁵ We check the sensitivity of our results to this assumption below.

We examine whether older analysts are more or less likely to deviate from the consensus than younger analysts as a test of the relationship between career concerns and herding.

Table 2 reports some descriptive statistics of these measures of herding from the I/B/E/S database. We will discuss these statistics in detail when we consider our findings in Section 4 below. Before then, we lay out our empirical strategy.

3. Empirical Strategy

This section discusses our strategy to measure the effect of analyst job tenure on the propensity to herd. First, we characterize the strengths and weaknesses of various methods of identifying the effect of job tenure on earnings forecast behavior. Then, we present the regression specifications we use to measure that effect.

3.1. Methods of Identifying the Effect of Job Tenure

There are several potential approaches to measuring how job tenure affects the propensity of a security analyst to herd with other analysts. The most direct method is to identify analysts of different ages and observe which analysts are more likely to herd. That is, if we know the age of security analysts, we can compare the performance of older analysts to younger analysts and characterize the difference between them as the effect of job tenure.

Although this approach is straightforward, there are several implicit assumptions that need to be made for this technique to accurately measure the effect of tenure on herding propensity. The most important is that older analysts must be exactly the same as younger analysts on average except for the fact that they are older. There must be no systematic reason why older analysts produce different earnings forecasts than younger analysts beyond the fact that older analysts have more years of job tenure.

This assumption might be particularly problematic for the security analyst labor market because job turnover is substantial. Analysts often leave our I/B/E/S sample before 1996; they might have found a better job, retired, or been fired by their investment bank. If the propensity of an analyst to leave the profession is related to her forecast ability, and forecast ability affects the propensity of herding, then attrition problems can substantially bias measures of herding by age using this technique. For example, if the best performing analysts often leave their positions to become mutual fund managers, then, on average, older analysts will be less skilled than younger analysts and the assumption that they only differ by job tenure will be violated. On the other hand, if investment banks tend to fire analysts for poor performance, then older analysts on average are more skilled than younger analysts and again comparisons in the performance of the two groups will be problematic.

Since analysts probably leave the I/B/E/S sample for both types of reasons, it is difficult to evaluate the size or direction of this attrition bias. However, there are other methods of gauging the effect of job tenure on herding that are less sensitive to this problem. One approach is to use the panel nature of the I/B/E/S database and observe how an individual analyst's propensity to herd changes over time as she gains experience. That is, compare the forecast performance of an analyst when she is young to her performance when she is older and characterize the difference as the effect of tenure.

This method avoids the attrition problem described above; the older and younger analysts being compared will have the same forecasting ability on average because they are the same analyst. However, there is another assumption that needs to be made for this procedure to produce useful estimates of how tenure affects the propensity to herd. The forecasting behavior of the analyst when she is older must be the same on average as her behavior when she is younger except for the effect of extra experience. No other factors that affect herding behavior must be changing over time for this analyst. This assumption could be questionable if there are aggregate changes over time in forecasting technology.

The I/B/E/S database does not contain any direct information about the age or experience of analysts that produce earnings forecasts. Therefore, we are unable to use the first approach of measuring the effect of tenure by comparing the herding measures of analysts of different ages. However, since the I/B/E/S is a long panel of analyst's forecasts over time, we can use the second method of identifying the tenure effect by simply observing analysts' forecast behavior as they age. As described above, our technique will lessen worries about certain attrition biases in the data, but we must be careful to consider the effects of aggregate changes in forecasting behavior.⁶ When we present the results of our work, we will attempt to gauge how these changes might be affecting our findings.

3.2. Regression Framework

Our empirical strategy is to relate measures of an analyst's herding propensity to the number of years that she has been an analyst, controlling for other factors that affect forecasting performance. We use regressions of the form:

$$\text{Herding Measure}_{i,j,t} = \alpha + \beta_1 \text{Tenure}_{i,t} + \text{Analyst} \times \text{Stock Dummies} + \text{Bank Dummies} + \varepsilon_{i,j,t} \quad (2)$$

where $\text{Herding Measure}_{i,j,t}$ is one of the two forecasting measures described above, $\text{Tenure}_{i,t}$ is our gauge of the job tenure of the analyst when she made the earnings forecast, $\text{Analyst} \times \text{Stock}$ is a full set of dummy variables for each analyst by stock cell, Bank is a set of effects for the investment bank that employs the analyst when the forecast is made, and $\varepsilon_{i,j,t}$ is an error term. Since we do not have any direct information on the years of experience of the analyst at the time she makes a forecast,

⁶ This second approach alleviates the attrition problem if good or poor analysts are on average more likely to leave the profession over time. However, there is still a concern with attrition bias if analysts whose performance is *improving* or *deteriorating* over time are leaving the sample.

$Tenure_{i,t}$ is a measure of the number of years we have observed the analyst in the I/B/E/S sample before she makes this earnings estimate.

This regression specification identifies the effect of job tenure on herding by comparing an analyst's forecasting behavior as she ages, the second empirical approach described in the previous section. Including analyst effects in the regression eliminates variation in tenure across analysts, leaving only within-analyst age variation (an analyst aging over time) to identify the tenure effect.⁷ Because of this fixed-effects strategy, we must only be concerned with factors that affect an individual analyst over time that might bias our estimates (such as changes in forecast technology over time) and not factors that influence some analysts but not others.

We also include stock effects in the specification to control for any differences in the difficulty of forecasting earnings of firms. One concern is that there might be differences across stocks in the tendency for analysts to depart from the consensus. If analysts that stay in our sample longer are more or less likely than other analysts to follow firms for which this propensity is greater, or they are more or less likely to move to cover those firms over time, then we could be measuring a spurious relationship between the age of the analyst and herding behavior. To control for these potential firm coverage differences, we include stock effects in the regression specification.

Since the same analyst often covers the same firm over many years, we can control for differences across firms in a more complete way by including analyst \times stock effects instead of just stock effects. In general, we obtain similar results if we include analyst and stock effects in the regression specification separately instead of using analyst \times stock effects. Hence, we only report the

⁷ One potential method of controlling for changes over time in forecasting technology is to also include year effects in the regression specification. However, since changes in an analyst's tenure are collinear with changes in calendar time, we are unable to identify the tenure effect including year effects. We account for trends over time in forecasting in other ways below.

regressions with analyst×stock effects and identify the effect of aging on forecast performance by following an analyst as she forecasts an individual firm over time.

A final potential problem with our regression specification is that analyst forecast performance may be affected by which investment banks hire them. Though unlikely, smaller and less respected investment banks might be more likely to encourage their analysts to herd than other investment banks. If analysts are moving to or away from these banks as they age, then we could be measuring a spurious relationship between tenure and the propensity to herd.⁸ We include investment bank effects to the model to control for such potential spurious relationships.

4. Results

In this section, we introduce the findings of our empirical analysis. First, we present our results of the effect of job tenure on the timing of forecasts. Then, we show our evidence for the effect of tenure on deviations from the consensus. Finally, we include checks of how sensitive our results are to how we measure herding behavior.

4.1. Timing of Forecasts and Job Tenure

We first examine how job tenure affects the probability that an analyst submits the initial earnings forecast for a firm. We estimate the regression model of Equation (2) using as a dependent variable an indicator that the analyst produced the first forecast for the firm. The coefficient on job tenure tells us whether security analysts with more experience are more or less likely to lead.

⁸ Analysts would have to be switching banks for this to be a problem. If an analyst does not move, then the analyst fixed effect in the regression specification would capture the effect of the investment bank.

Table 3 displays the results of these regressions. Column (1) presents a regression specification that includes only analyst \times stock effects.⁹ Older security analysts are more likely to produce the first earnings forecast of a stock than younger analysts. The coefficient on age is statistically significant from zero and implies that ten years of additional tenure raises the probability of leading by about 3.8 percentage points. Since the probability that analysts in this I/B/E/S sample produced the first forecast for a firm was 7.7%, this implies that 10 years of aging increases the propensity to lead by 50%.¹⁰ In column (2) of Table 3, we add the investment bank effects with the analyst \times stock effects. The effect of experience on the probability of producing the first forecast increases, but the difference is not qualitatively important.

There are additional complications involving which investment bank an analyst works for. A firm that an analyst follows might be a client of the analyst's investment bank. Such relationships have been documented to bias forecasts, though not the herding measures we are considering *per se* (see Michaely and Womack (1996)). In our context, we worry that to placate a client, an analyst might produce very early forecasts for the firm. If analysts are moving to investment banks over time that have as clients the firms they cover, then our herding findings might be an indication of this movement. One method of controlling for this potential problem is adding analyst \times stock \times bank effects to the regression specification.¹¹ Column (3) of Table 3 displays the estimates of the effect of tenure on probability of leading including these dummies in the model. Again, the effect of tenure is not

⁹ The standard errors of all the regressions presented are calculated to allow for intra-analyst correlation in the error term.

¹⁰ We only present the results of linear probability models. However, results from logit and probit models are qualitatively similar.

¹¹ Since not all firms are clients of an investment bank, adding just investment bank dummies does not solve this problem.

significantly different from our previous estimates, implying that client relationships are not driving our timing results.

We view this timing evidence as providing substantial support for a link between career concerns and herding for several reasons. First, it is difficult to argue that our timing results are being driven by aggregate changes in forecasting technology. Such changes over time might affect when analysts make their first forecast, but they would probably influence more and less experienced analysts equally. For example, if it is becoming more difficult to forecast earnings over time, then analysts might decide to delay their initial forecasts later than before so that they can receive more information. But analysts of all ages would do that, so the order of forecasts would not be affected.

Second, analysts have discretion over when to generate forecasts. In particular, the timing of their forecasts are not artifacts of investment banks telling older analysts to forecast first and younger analysts to forecast later since investment banks generally have one analyst covering a particular stock within an industry.

Finally, generating timely forecasts is highly valued in the investment community and one of the benchmarks with which analysts are often evaluated in surveys. In general, generating timely forecasts requires more effort than delaying one's forecasts and waiting for the consensus to form and then merely submitting the consensus forecast as one's own (see, e.g., Stickel (1992)). Because of the effort involved, this timing finding (as we will argue below) goes a long way in ruling out a subset of reputation/herding stories and also some alternative non-reputation/herding stories. This point becomes an especially important fact to keep in mind when interpreting the other findings below.

4.2. Deviation from Consensus and Job Tenure

If we view the timing result as supportive of older analysts herding less, then we should expect to find that more experienced analysts also deviate from the consensus forecast more than their younger

counterparts. Otherwise, interpreting our first finding becomes problematic. In Table 4, we present the results of the regression specified in Equation (2) with the *Deviation from Consensus* measure of herding defined in Equation (1) as the dependent variable.

Column (1) of Table 4 reports the results of this regression including only analyst×stock effects in the specification. The coefficient on years of experience suggests that analysts deviate more from the consensus earnings forecast as they age. The magnitude of the coefficient implies that 10 extra years of experience increases an analyst's deviation from consensus forecasts about 3.7 cents, or about one-third of the difference between the 75th percentile and the median of the consensus deviation measure for the entire sample. (See Table 2 for the summary statistics on the distribution of the consensus deviation measure.) Adding investment bank effects in column (2) of Table 4 and controlling for client relationships by including analyst×stock×bank effects in column (3) does not substantially change the estimate on tenure.

Since we found above that more experienced analysts produce earlier forecasts, we need to investigate whether older analysts' forecasts deviate more from the consensus than younger analysts merely because they submit earlier forecasts on average. Since more information is available to those who submit forecasts later rather than earlier, the dispersion in forecasts may simply reflect the fact that older analysts had less information when they submitted their forecasts. We can control for this problem by adding month effects in our regression specification. That is, we add to the right hand side of our regression specification a set of dummies for the month in which a forecast is submitted in a given year.¹² Column (4) of Table 4 displays the results of a regression specification that includes month effects with analyst×stock effects. Again, the coefficient on experience does not qualitatively change, suggesting that our timing finding is not driving our deviation from consensus results.

Having established this second finding, it becomes more difficult to simultaneously rationalize both the timing result and the deviation from consensus result with arbitrary stories unrelated to reputation/herding like stories.¹³

4.3 Sensitivity Analysis

We now perform a sensitivity analysis of our findings to rule out uninteresting, mechanical explanations that might be hard-wiring our results. With regard to the first finding on timing of forecasts, we primarily have to be careful about how we construct our sample. Recall from our discussion in Section 2.1 that we want to only include analysts who had an opportunity to make the initial estimate for a firm in a given year. We did so by arbitrarily limiting the sample to analysts who have followed a stock for at least four years. We find that our results are unaffected by whether we use less than or more than four years as a cutoff. For brevity, we do not show this result and focus our discussions instead on the sensitivity of our second finding on deviation from the consensus forecast, since this result is more prone to potential biases.

One of the most serious issues regarding our methodology for examining the relationship between job tenure and analysts' deviations from consensus forecasts involves the possibility that our results are being driven by time trends in forecast behavior. If all analysts, regardless of experience, are deviating more from the consensus over time for some reason, then our regression results would

¹² Since forecasts are only included in the sample if they were made between January 1st and July 1st of the year studied, the month dummies represent the months January through June.

¹³ One could of course create behavioral stories that might rationalize the findings. One such story often told to us is that older analysts trust their information more than younger ones. Such stories are not convincing since idiosyncratic behavioral biases do not necessarily imply copying the actions of other analysts and it isn't entirely clear how they translate into timing or dispersion. Regardless, the point of herding models is that rational agents may discard their own information to mimic the actions of others. Admittedly, others more creative than us can come up with stories; however, our view is the reputation/herding mechanism is the more parsimonious one.

spuriously indicate that analysts are less likely to herd as they age. We now attempt to gauge the extent of this potential problem. Then we examine how sensitive our deviation from consensus results are to how we defined the window for calculating analysts' earnings forecasts.

If firms are becoming more difficult to forecast over time, or analysts are deviating from consensus forecasts more over time for other reasons, then the standard deviation of analysts' earnings forecasts for a firm should be increasing over time also. For each firm that an analyst covers in a year, we calculate the standard deviation of the earnings forecasts of all the other analysts that cover that firm: $\sigma(F_{-i,j,t})$. We then include this standard deviation measure as an independent variable in the regression specification of Equation (2) to control for the changes over time in the dispersion of other analysts' forecast for the firm (see Table 2 for the descriptive statistics of the standard deviation measure). Column (1) of Table 5 presents the results of this regression including analyst \times stock effects in the specification. The coefficient on the standard deviation variable is the expected sign; when the forecasts of other analysts are more disperse, an analyst's forecast is likely to be farther away from the consensus forecast. But including the standard deviation measure does not affect the estimated effect of tenure on deviating from the consensus, suggesting that time trends are not driving our results.

Next, we examine whether some one-time technological shock in forecasting is driving our findings. For example, if a shock makes forecasting very difficult for analysts late in the sample period, causing forecasts to deviate more from the consensus, then we would spuriously find that older analysts are less likely to herd. Column (2) of Table 5 reports the deviation from consensus regression result using the first seven years of the I/B/E/S sample (1983-1989); whereas, column (3) presents the estimates for the final seven years (1990-1996). Although the coefficient on job tenure is higher for the earlier sample, both estimates suggest a large positive effect of experience on the propensity to deviate

from consensus forecasts. Therefore, our results do not appear to be driven by a sudden aggregate change in forecasting difficulty.

Having established that our results do not appear to be driven by time trends, we finally examine how sensitive our deviation from consensus results are to a couple of decisions we make in constructing the herding measure. First, we change the cutoff date (i.e. July 1st). In column (4) of Table 5, we present the regression result examining how job tenure affects an analyst's deviation from the consensus forecast when October 1st is used as the cutoff date instead. The sample size increases because more analysts produce earnings forecasts for a firm by October 1st than July 1st of the forecast year. However, the coefficient on the effect of experience does not change from the similar regression with July 1st as the cutoff. Column (5) displays the results where March 1st is the cutoff date. Again, there is a positive effect of job tenure on deviations of forecasts from the consensus.

In our previous work, only ten analysts must be covering a firm for it to be included in the sample. Now we examine how our results change if we exclude firms with fewer than twenty analysts covering it. Column (6) of Table 5 presents the estimates of the effect of job tenure on deviation from the consensus for this restricted sample of firms. The coefficient on job tenure is slightly larger than the estimate that includes firms with sparser coverage in the sample; however, the two coefficients are not substantially different.¹⁴

To create our measures of herding, we needed to make some decisions about which earnings estimates to consider and which firms to include in our sample. Our robustness checks suggest that our choices do not drive our findings that younger security analysts are more likely to herd than older analysts. Also, our evidence suggests that these results are not caused by aggregate time trends in forecasting behavior.

5. Auxiliary Results

Up to this point, we have focused primarily on one aspect of recent models of reputation and herding, the link between career concerns as measured by job tenure and various measures of herding. We next consider two other forecast performance measures: forecast accuracy and frequency of revisions of earnings forecasts. We find some evidence for variation of these measures with job tenure. However, we view these additional findings as auxiliary results because existing reputation-based herding theories have less direct implications for them. However, it is still informative to relate these two measures to the main findings to enrich our picture of the mechanisms driving the submission of earnings forecast.

5.1. Forecast Accuracy and Job Tenure

First, we consider the relationship between forecast accuracy and job tenure. It is well known that earnings forecasts submitted closer to are more accurate than those submitted farther from the actual earnings report date.¹⁵ For a firm that ends its fiscal year in December, an estimate made in August will likely be more accurate than an estimate made the previous January. Therefore, we need a method for determining which earnings forecast of an analyst to use in our analysis to make our accuracy measure comparable across analysts. We adopt the same rules for considering an earnings forecast as we used in constructing the deviation from consensus measure. The latest forecast between January and July of the fiscal year being considered by the analyst is used as the analyst's forecast for

¹⁴ Not surprisingly, our results are also qualitatively unchanged if we expanded our sample size by requiring that only five analysts must be covering a firm for it to be included in the sample.

¹⁵ See Crichfield, Dyckman and Lakonishok (1978).

that firm. A natural measure of analyst accuracy is to compare this forecast and the actual earnings of the firm that fiscal year:

$$\text{Forecast Accuracy}_{i,j,t} = |F_{i,j,t} - A_{j,t}| \quad (3)$$

where $F_{i,j,t}$ is again the earnings per share (EPS) estimate of the analyst and $A_{j,t}$ is the actual EPS.

The regression results for the effect of age on this forecasting accuracy measure are reported in Table 6. The regression specification is again Equation (2) with the dependent variable the forecasting accuracy measure. A positive coefficient on the job tenure variable suggests that older analysts are less accurate than younger ones. The coefficient on experience in column (1), which includes only analyst×stock effects in the specification, indicates that an analyst’s accuracy deteriorates as she ages. Ten years of aging decreases the forecasting accuracy measure by about 13 cents, or about 40% of the difference between the 75th percentile and the median of the forecast accuracy measure (see Table 2 for descriptive statistics). The results including the investment bank dummies are shown in column (2). The coefficient on age is again positive and slightly larger than column (1). Adding analyst×stock×bank effects in column (3) even strengthens the conclusion that an analyst forecasts less accurately as she ages.¹⁶

However, these accuracy results might be difficult to interpret because of attrition problems in the I/B/E/S sample. For example, if analysts are fired because of poor accuracy performance, then, because of this selection, we would expect to only observe analyst performance deteriorating over time. Analysts who started poorly were fired and not given the opportunity to improve. Only the analysts that started well stayed in the sample and we would observe the performance of some of those

¹⁶ Adding month effects to the regression specification also does not affect the coefficient on job tenure.

survivors deteriorate. If this kind of attrition is important, then it would be difficult for us to find that more experienced analysts are more accurate.

Therefore, we view these accuracy results as only suggestive, probably only ruling out large learning by doing (or experience) effects in the production of forecasts. One might quickly jump to the conclusion that analysts' forecast accuracy is deteriorating with more experience because analysts might be getting lazier as they age. One formal story consistent with this is Holmstrom's (1982) seminal career concern model in which agents exert less effort over time. However, such a conclusion is unconvincing in this instance because the easiest way to slack in this labor market is to delay one's forecast, wait for the consensus to develop, and free ride off the information of others. In fact, we see the opposite. Older analysts issue more timely and bolder forecasts than younger analysts. Rather, the most cautious interpretation of the accuracy finding is that it does not radically add to nor change the conclusion that there is some reputation/herding mechanism at work in the submission of earnings forecasts.¹⁷

5.2. Frequency of Revisions and Job Tenure

Finally, we consider the relationship between frequency of revisions of earnings forecasts by an analyst and job tenure. We can construct a revision measure by considering the number of times an analyst revises her earnings forecast of a company in a year and comparing that to how often other analysts covering the same stock in the year revise their estimates. Specifically, we count the number

¹⁷ To be complete, it is interesting to note that Scharfstein and Stein (1990) can conceivably be extended to simultaneously account for the two herding as well as the accuracy results. In their canonical model, agents are judged according to absolute performance and relative action. They point out however that relative ability concerns may offset herding tendencies due to a "superstars" effect. So if relative performance matters for analysts, then older analysts may refrain from taking *ex-ante* accurate, unbiased forecasts and may end up taking bolder but less accurate forecasts to stand out.

of times that an analyst revises her earnings estimate between January 1st and July 1st and compare that to the average number of revisions of other analysts covering that stock over the same time period:

$$\text{Revision Performance}_{i,j,t} = R_{i,j,t} - \frac{1}{n} \sum_{m \in -i} R_{-i,j,t} \quad (4)$$

where $R_{i,j,t}$ is the number of times analyst i revises her earnings estimate of stock j in year t during the fiscal year before July 1st. We subtract by the average number of times that other analysts revise their forecasts to control for heterogeneity across stocks in the propensity of analysts to revise earnings estimates.

We estimate the regression model of Equation (2) using as a dependent variable now the number of revisions that an analyst made less the average number of revisions made by others following the stock in the same fiscal year. The coefficient on job tenure measures whether older analysts revise their estimates more often or less often than younger analysts.

Table 7 presents the estimates for the regression model with the revision performance dependent variable. The specification in column (1) includes only analyst×stock effects. The coefficient on age is negative and statistically significant from zero, suggesting that more experienced analysts are less likely to update their forecasts. The magnitude of the coefficient on age in column (1) implies that aging ten years lowers our revision performance measure by 0.104, or about 20 percent of the difference between the 75th and 50th percentile of the distribution of the revision measure for the entire sample (see Table 2). Including investment bank effects in column (2) of Table 7 or analyst×stock×bank effects in column (3) does not alter the coefficient on age qualitatively.

One potential problem with these regressions is that all analysts might be revising their forecasts less over time. However, since we are adjusting the revision measure by the average number of revisions of other analysts covering the stock, we are controlling for changes in these aggregate

propensities to revise forecasts. The frequency of revision result is an additional fact that may help inform future theories that purport to explain the mechanisms behind the generation of earnings forecasts.

6. Conclusion

In this paper, we document a link between career concerns and herding in the labor market for security analysts. Using job tenure as a proxy for career concerns, we find that more experienced analysts are more likely to issue timely earnings forecasts for a firm and that they also deviate more from the consensus forecast than younger analysts. We argue that these two findings suggest that reputation mechanisms are at work behind the submission of forecasts. Furthermore, we find weaker evidence that older analysts' forecasts are not better than younger analysts and older analysts revise their forecasts less frequently than their younger counterparts.

At least two questions for future research arise out of this paper. First, does the market for earnings forecasts recognize the link between career concerns and herding among security analysts? One could answer this question by looking at whether a firm's stock price reacts differently to forecasts generated by older versus younger analysts. Such a study is interesting as it would complement the findings in this paper as well as speak to the efficiency of capital markets. Second, do analysts herd in other aspects of earnings forecasts? We have largely focused on only a few herding measures related to forecast performance. There are others. For instance, Abarbanell and Bernard (1993) document that analysts' forecasts under-react to the information in recent earnings. It would be interesting to study to what extent this under-reaction varies with the job tenure of analysts. There are undoubtedly other questions. We leave them for future research.

References

- Abarbanell, Jeffrey S. and Victor C. Bernard (1992) "Tests of Analysts' Overreaction/Underreaction to Earnings Information as an Explanation for Anomalous Stock Price Behavior," *Journal of Finance*, 47, 1181-1207.
- Banerjee, Abhijit (1992) "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107, 797-817.
- Bikhchandani, Sushil, David Hirshleifer and Ivo Welch (1992) "A Theory of Fads, Fashions, Custom and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100, 992-1026.
- Chevalier, Judith A. and Glenn D. Ellison (1997) "Career Concerns of Mutual Fund Managers," MIT Working Paper.
- Crichfield, Timothy, Thomas R. Dyckman and Josef Lakonishok (1978) "An Evaluation of Security Analysts' Forecasts," *Accounting Review*, vol. 53, no. 3, 651-668.
- Ehrbeck, Tilman and Robert Waldmann (1996) "Why are Professional Forecasters Biased? Agency versus Behavioral Explanations," *Quarterly Journal of Economics*, 111, 21-40.
- Falkenstein, Eric (1996) "Preferences for Stock Characteristics as Revealed by Mutual Fund Portfolio Holdings," *Journal of Finance*, 51, 111-135.
- Fama, Eugene F. (1980) "Agency Problems and the Theory of the Firm," *Journal of Political Economy*, 88, 288-307.
- Graham, John R. (1998) "Herding among Investment Newsletters: Theory and Evidence," *Journal of Finance*, forthcoming.
- Grinblatt, Mark, Sheridan Titman and Russ Wermers (1995) "Momentum Investment Strategies, Portfolio Performance, and Herding: A Study of Mutual Fund Behavior," *American Economic Review*, 85, 1088-1105.
- Holmstrom, Bengt (1982) "Managerial Incentive Problems: A Dynamic Perspective" in *Essays in Economics and Management in Honor of Lars Wahlbeck*, Helsinki: Swedish School of Economics, 209-230.
- Huddart, Steven (1996) "Reputation and Performance Fee Effects on Portfolio Choice by Investment Advisors," Working Paper, Duke University.
- Lamont, Owen (1995) "Macroeconomic Forecasts And Microeconomic Forecasters," NBER Working Paper 5284.
- Lazear, Edward and Sherwin Rosen (1981) "Rank-Order Tournaments as Optimum Labor Contracts," *Journal of Political Economy*, 89, 841-64.

- Lakonishok, Josef, Andrei Shleifer, Richard Thaler and Robert Vishny (1991) "Window Dressing by Pension Fund Managers," *American Economic Review*, 81, 226-231.
- Michaely, Roni and Kent Womack (1996) "Conflict of Interest and the Credibility of Underwriter Analyst Recommendations," Working Paper, Cornell University.
- Nocera, Joseph (1997) "Who Really Moves The Market?" *Fortune*, October 27, 90-110.
- Nofsinger, John and Richard Sias (1996) "Herding by Institutional and Individual Investors," Working Paper, Marquette University.
- Peles, Nadav (1993) "The Determinants of Institutional Demand for Equity: An Empirical Study," Working Paper, Columbia University.
- Prendergast, Canice and Lars Stole (1996) "Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning," *Journal of Political Economy*, 104, 1105-34.
- Scharfstein, David S. and Jeremy C. Stein (1990) "Herd Behavior and Investment," *American Economic Review*, 80, 465-479.
- Stickel, Scott E. (1992) "Reputation and Performance Among Security Analysts," *Journal of Finance*, 47, 1811-1836.
- Trueman, Brett (1994) "Analyst Forecasts and Herding Behavior," *The Review of Financial Studies*, 7, 97-124.
- Welch, Ivo (1996) "Herding among Security Analysts," UCLA Finance Working Paper #8-96.
- Wermers, Russ (1995) "Herding, Trade Reversals and Cascading by Institutional Investors," Working Paper, University of Colorado, Boulder.
- Wylie, Samuel (1996) "Tests of The Accuracy of Measures of Herding using UK Data," Working Paper, London Business School.
- Zwiebel, Jeffrey (1995) "Corporate Conservatism and Relative Compensation," *Journal of Political Economy*, 103, 1-25.

Table 1: Descriptive Statistics of I/B/E/S Sample

	Mean	10 th Percentile	90 th Percentile
	(1)	(2)	(3)
Number of Years Analyst is in Sample	4.05 [3.70]	1	10
Number of Analysts that Cover a Firm in a Year	18.59 [10.57]	7	34
Number of Stocks Analyst Follows in a Year	9.43 [4.71]	4.5	15.47
Number of Years an Analyst Follows a Stock	2.73 [2.63]	1	6

Notes: Data on 8421 analysts are from I/B/E/S. The sample period is from 1983 to 1996. Standard deviations are in brackets.

Table 2: Descriptive Statistics of Measures of Herding

	Mean	25 th Percentile	50 th Percentile	75 th Percentile
	(1)	(2)	(3)	(4)
Probability that an Analyst Produces the Initial Forecast of a Firm's Earnings	0.077			
Deviation from Consensus: $ F_{i,j,t} - \bar{F}_{-i,j,t} $	0.20 [0.89]	0.02	0.06	0.17
Standard Deviation of Other Analysts' Forecasts of Firm: $\sigma(F_{-i,j,t})$	0.25 [0.85]	0.05	0.11	0.25
Additional Forecast Performance Measures:				
Forecast Accuracy: $ F_{i,j,t} - A_{j,t} $	0.82 [8.24]	0.05	0.16	0.50
Revision Performance: $R_{i,j,t} - \bar{R}_{-i,j,t}$	-0.00 [0.96]	-0.64	-0.22	0.50

Notes: Data from I/B/E/S for the years 1983 through 1996. 198,173 observations. Standard deviations are in brackets. See the text for a description of the sample used to calculate the probability that an analyst produces the leading forecast.

**Table 3: The Effect of Tenure on the Probability
that an Analyst Produces the Earliest Forecast**

	(1)	(2)	(3)
Tenure of Analyst	.0038 (.0016)	.0045 (.0016)	.0055 (.0018)
Analyst× Stock Effects	Yes	Yes	---
Investment Bank Effects	No	Yes	---
Analyst× Stock× Investment Bank Effects	No	No	Yes
Observations	104,807	104,807	104,807

Notes: The sample consists of analysts with at least four years of experience covering a given firm in the I/B/E/S dataset from the years 1987 through 1996. The dependent variable is an indicator that the analyst was the first to produce a forecast of the earnings for this stock. The average probability that an analyst in this sample produces the first forecast is 7.7%. Robust standard errors are in parenthesis. They are calculated allowing for intra-analyst correlation.

**Table 4: The Effect of Tenure on the
Deviation of Analysts' Forecasts from the Consensus**

	(1)	(2)	(3)	(4)
Tenure of Analyst	.0037 (.0013)	.0037 (.0013)	.0056 (.0015)	.0038 (.0012)
Analyst× Stock Effects	Yes	Yes	---	Yes
Investment Bank Effects	No	Yes	---	No
Analyst× Stock× Investment Bank Effects	No	No	Yes	No
Month Effects	No	No	No	Yes
Observations	198,173	198,173	198,173	198,173

Notes: Standard errors are in parenthesis. They are calculated allowing for intra-analyst correlation. The dependent variable is a measure of the deviation of an analyst's forecast from the consensus forecast as described in the text.

Table 5: The Effect of Tenure on the Deviation of Analysts' Forecasts from the Consensus, Sensitivity Tests

	Standard Deviation	1983-1989	1990-1996	October 1 st Cutoff	March 1 st Cutoff	At Least 20 Analysts
	(1)	(2)	(3)	(4)	(5)	(6)
Tenure of Analyst	.0035 (.0008)	.0173 (.0038)	.0069 (.0024)	.0054 (.0014)	.0051 (.0017)	.0055 (.0018)
Standard Deviation of Other Analysts' Forecasts of Firm: $\sigma(F_{-i,j,t})$.8383 (.0763)					
Analyst \times Stock Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	198,173	103,526	94,647	262,250	66,536	95,550

Notes: Standard errors are in parenthesis. They are calculated allowing for intra-analyst correlation. The dependent variable is a measure of the deviation of an analyst's forecast from the consensus forecast as described in the text.

Table 6: The Effect of Tenure on the Forecast Error

	(1)	(2)	(3)
Tenure of Analyst	.0133 (.0069)	.0152 (.0076)	.0205 (.0091)
Analyst× Stock Effects	Yes	Yes	---
Investment Bank Effects	No	Yes	---
Analyst× Stock× Investment Bank Effects	No	No	Yes
Observations	198,173	198,173	198,173

Notes: Standard errors are in parenthesis. They are calculated allowing for intra-analyst correlation. The dependent variable is the forecast accuracy measure described in the text.

**Table 7: The Effect of Tenure on the
Number of Times Analysts Revise an Earnings Forecast**

	(1)	(2)	(3)
Tenure of Analyst	-.0093 (.0028)	-.0089 (.0029)	-.0123 (.0035)
Analyst× Stock Effects	Yes	Yes	---
Investment Bank Effects	No	Yes	---
Analyst× Stock× Investment Bank Effects	No	No	Yes
Observations	198,173	198,173	198,173

Notes: Standard errors are in parenthesis. They are calculated allowing for intra-analyst correlation. The dependent variable is the difference in the number of times that an analyst revises an earnings estimate, and the number of times analysts revise the forecast for that stock on average.