

A Catering Theory of Analyst Bias

RICHARD LAI

Abstract

We posit a theory that runs counter to how conventional wisdom thinks about analyst bias, that it is the result of distorted incentives by “upstream” factors like the analysts’ employers. We suggest that analysts are also heavily influenced by the beliefs of investors downstream, the purported victims of analyst bias. We adapt Mullainathan-Shleifer’s theory of media bias to build a theory of how analysts cater to what investors believe. The theory also predicts that competition among analysts does not reduce their bias. We provide empirical support for this theory, using an enormous dataset built from over 6.5 million analyst estimates and 42.8 million observations on investor holdings, which we argue is a proxy for what investors’ beliefs. We use a simultaneous-equations model for estimation, with instruments to rule out alternative interpretations of the direction of causality. For additional robustness, we investigate the time series of analyst bias and heterogeneity in investor beliefs from 1987 through 2003. Dickey-Fuller tests show that both have unit roots, but we establish that cointegration holds. Further, we employ a vector-autoregressive model to show Granger-causality between the two.

The bias of analyst forecasts has received extensive empirical study, from as early as Mastrapasqua and Bolten (1973) and Brown and Rozeff (1978). A recent survey is in O’Brien (2003). A large class of studies deals with how analysts’ biases might affect stock prices and other aspects of the market. This paper joins an emerging class of studies that seeks to understand the source of the bias. For example, Hong and Kubik (2003) show that analysts who are positively biased get better jobs in the industry. Cowen, Groysberg and Healy (2003) provide evidence that analysts whose employers are more commission-based, such as pure brokerages, are more bias. Underlying many of these explanations is an *upstream* source of analyst bias: that of distorted incentives by their employers and the companies covered.

In this paper, we ask the seemingly obvious question: might not analysts’ biases be a result of their catering to *downstream* influence - that is, their direct clients, the readers? This paper can also be viewed as a contribution to the already large body of literature on behavioral corporate finance, in which a theme is that firm behavior can be explained by their responses to clients and the market. This theme has been used to explain capital structure (Baker and Wurgler (2002)), dividend payouts (Baker and Wurgler (2003)), and acquisitions (Shleifer and Vishny (2003)).

Our “catering” theory, if true, has important practical implications. For example, it means that policy makers have to go beyond reforming incentives and organizational structures inside analysts’ institutions, if we were to have any hope of reducing analyst

bias. Finally, this research question opens up a range of others. The most important is how readers of analyst reports are in turn influenced by what analysts write.

We start with Section I, in which we outline a model of how analysts might cater to their readers' beliefs. We adapt the model of Mullainathan and Shleifer (2003) (hereafter called "M-S") for media bias for our purpose. In Section II, we develop the propositions resulting from the theory. The key propositions are that that heterogeneity of investor beliefs increases analyst bias, and surprisingly, so does competition among analysts. In Section III, we describe how we measure the variables in the theory, and provide detailed descriptions of the construction of our datasets, using over 6.5 million analyst estimates from FirstCall, 42.8 million observations on investor holdings from Thomson Financial, and other databases such as Hoover's Online. In Section IV, we show the empirical results of our estimation. We employ a simultaneous-equations model for estimation, with instruments to rule out alternative interpretations of the direction of causality. For additional robustness, we investigate the time series of analyst bias and heterogeneity in investor beliefs from 1987 through 2003. Dickey-Fuller tests show that both have unit roots, but we establish that cointegration holds. Further, we employ a vector-autoregressive model to show Granger-causality between the two. In Section V, we conclude with some thoughts on what questions open up for research given the results of this paper.

I. A Model

Our model is almost wholly adapted from Mullainathan and Shleifer (2003) and the well-known models for competition pioneered by Hotelling. We see our contribution in the empirics. To make the model intuitive, our setting involves analysts who forecast financial results, such as a stock's earnings per share. These forecasts are for institutional investors, since few analysts write for retail clients. Examples of institutional investors might include pension funds, university endowments, mutual funds, and private equity funds. The algebra follows.

A. Analysts

Analysts interact with their clients, the institutional investors. In particular, the analysts find out investors' beliefs about the stock she is covering, if any. This can be done by checking public information on investors' holdings, or through frequent interactions with them.

To write their reports for investors, analysts visit companies whose stocks they cover, attend company announcements, or participate in conference calls. Analysts covering the same stock are assumed to receive the same information; companies do not make proprietary disclosure to analysts, as is the case in practice.

For concreteness and without loss of generality, let us say analysts get information on a stock's future earnings per share d , which is distributed $N(t, \sigma)$.

On having an idea of investors' beliefs, analysts choose their slant $s(d)$. If two analysts cover the same stock, they choose s simultaneously. Then they get data from companies, which is some noisy measure of d . (The noise is null when, for example, companies finally announce their actual earnings per share.) After getting their noisy data, the analysts write up reports $n = d + s$ for their investors.

B. Investors

Investors are interested in d . We examine two scenarios of investors: (1) when they are rational, and (2) when they are “behavioral.”

Rational investors are unbiased and dislike slant by analysts. Their utility may be written as:

$$U = \alpha - \beta \cdot s(d)^2 \dots \dots \dots (1)$$

We assume that investors pay for analysts’ opinions indirectly, for the business the former gives to the latter now and then. In our setting, the business might be the brokerage business that investors channel to the analysts’ firms.

Behavioral investors, unlike rational investors, hold beliefs about d that may be biased¹, by an amount b . We assume that their beliefs are distributed $N(b, \sigma)$. Furthermore, they do not like getting analyst opinions different from their beliefs. The concept of “cognitive dissonance,” pioneered by Festinger (1957) and recently reviewed in Kahneman (2003), is consistent with this assumption. As an example, once investors have bought some stocks, they prefer opinions confirming that they have made the right purchases. Biased investors have utility of the form:

$$\begin{aligned} U &= \alpha - \beta \cdot s(d)^2 - \gamma(n - b)^2 \\ &= \alpha - \beta \cdot s(d)^2 - \gamma(d + s(d) - b)^2 \dots \dots \dots (2) \end{aligned}$$

Investors, rational or behavioral, read only analysts’ reports that generate positive expected utility. Analysts, of course, want their investors to read their reports, as that is correlated with the amount of business generated for their employers.

C. Cases Considered

We consider two simple cases: (1) where all investors are rational, and (2) where they are all behavioral.

In the behavioral case, we also consider two sub-cases: (2a) homogenous investors, a situation in which all investors hold the same beliefs b , such as when everyone is enthusiastic about Internet stocks, or (2b) heterogeneous investors, in which two investors $i \in [1, 2]$ hold beliefs that are uniformly distributed in $[b_1, b_2]$, $b_2 > 0$.

We examine two types of analyst industry structures: (1) a monopolist analyst as the only one covering a stock, and (2) duopolist analysts who cover the same stock. Of course, there can be different structures for different stocks. Also, we use the duopolist structure to illustrate “competition.”

Finally, we define average slant (AS) for each stock as:

$$AS = \begin{cases} E_d [(n - d)^2] & , \text{ homogenous} \\ \int_i E_d [(n_i - d)^2] & , \text{ heterogeneous} \dots \dots \dots (3) \end{cases}$$

where in the heterogeneous case, n_i is the report investor gets from analyst i , $i \in [1, 2]$.

¹ To keep the terms straight, notice that we have used the word “slant” for what analysts report and “bias” for what investors believe.

II. Analysis of the Model

We will go through the rational, homogenous, and then heterogeneous cases, and analyze what happens to slants with monopolist and duopolist industry structures for each case.

A. Rational Investors

This is the easy case: whether monopoly or duopoly, analysts cater to investors who dislike slant and so do not slant. This is straight-forward from the first-order conditions of (1).

In the duopoly case, the competition would be Bertrand-like, so that the implicit value of analysts' reports to investors is zero. To see this, if one analyst thinks he delivers value that can be appropriated, he might ask investors for more business. Investors will simply switch to another analyst, since all analysts provide the same unslanted reporting.

We summarize this case with the following proposition for empirical testing:

Proposition 1 – With rational investors, the average slant (AS) is zero, for both monopoly and duopoly industrial structures.

B. Behavioral Homogenous Investors

Investors might be homogenous when they are all caught up in a market run-up or down-turn. Homogeneity might also be the result of herding, such as the type Grinblatt, Titman and Wermers (1995) shows for mutual funds.

For the case of a monopolist analyst, since utility functions are separable in d , the first-order conditions of (2) give the optimal slant as:

$$s^*(d) = \gamma(b - d)/(\gamma + \beta) \dots\dots\dots(4)$$

For the anxious reader, table 1 at the end of this section summarizes how the amount of slant here compares with that in other cases.

As M-S noted, equation (4) provides a "theory of spin." Once a company provides a positive spin d on itself, analysts' biases are automatically derived from the homogeneity of investors, such as during an Internet bubble. With a dynamic model where investors' priors are shaped by previous analyst reports, this becomes a vicious cycle.

We also note these comparative statics:

$$\begin{aligned} \partial |s^*(d)| / \partial \gamma &> 0 \\ \partial |s^*(d)| / \partial \beta &< 0. \end{aligned}$$

The first says that the more investors hate dissonance, the greater the slant. The second says that the more investors hate slants, the smaller the slant.

With duopolist analysts, equation (4) holds again on their equilibrium path. As with the rational investors case, Bertrand-like competition holds.

Since the slants are the same in both monopolist and duopolist cases, the average slant AS for both are the same.

Proposition 2 – With behavioral homogenous investors, the average slant (AS) is $\gamma(b - d)/(\gamma + \beta)$, and is the same for both monopoly and duopoly industrial structures.

This is a surprising result, since it suggests that competition among analysts does not reduce slant. On reflection, however, it might even be intuitive. In cases where investor beliefs are somewhat homogenous, such as during the Internet bubble or when there are few investors all holding the same stocks (i.e., same imputed beliefs about the potential of these stocks), it is appealing to think that competition among analysts does not reduce their desire to cater to investors. For example, Graham (1999) provides evidence that investment newsletter writers tend to herd when the public signal (in our case, the uniform belief by many investors) is strong, even when the analysts' private (unslanted) information is inconsistent with this strong signal.

One counter-argument against this proposition is that with competition, might not an analyst do better by standing out with a different forecast from the herd? It turns out that such is not always the case. Indeed, it is easy to have payoff matrices such that enough other analysts deviate from the herd prediction to make "standing out" a poor and out-of-equilibrium strategy. What is desirable seems to be an empirical question, which we test.

C. Behavioral Heterogeneous Investors

For the monopolist analyst, the optimal slant is now:

$$s^*(d) = \gamma[(b_1 + b_2)/2 - d]/(\gamma + \beta) \dots\dots\dots(5)$$

Compared with the homogenous investors case, the amount of slant is greater if the average investor bias (the term $(b_1 + b_2)/2$ in equation 5) is larger than that in the homogenous case (b in equation 4).

Proposition 3 – When a stock is covered by a monopolist analyst, the average slant by the analyst is a function of the mean bias among the investors, whether for the homogenous or heterogeneous cases.

To prove equation (5), we first show that the optimal s^* has the form $\gamma(B - d)/(\gamma + \beta)$ and then ask what B will optimize s .

Suppose s^* has the desired form. Define $X = \{b_i \mid Eu_i(s(d)) = 0\}$ as the biases of investors receiving no utility. X is non-empty, because if it were, the monopolist analyst can slant a bit more to cater to more investors until the marginal investor is again receiving zero utility, thus violating the optimality of s^* .

Now X is a subset of $[b_1, b_2]$, so it has a well-defined *inf* and *sup*, which we will denote x_{inf} and x_{sup} , with the investors' corresponding utilities denoted u_{inf} and u_{sup} . Using Lemma 1 from the Appendix, an analyst strategy yields the maximin payoff given x_{inf} and x_{sup} and has the s^* form needed in this proof. Specifically, B is a linear combination of x_{inf} and x_{sup} . Further, u_{inf} and u_{sup} are positive, so that all investors satisfy their participation constraints.

The second and final step is to determine the specific B . This turns out to be easy. Since we are restricted to strategies of the form s^* where expected utilities are quadratic

in B , we are in a standard Hotelling model with quadratic transportation costs. Therefore, B must be the average of b_2 and b_1 .

In the duopolist analyst case, analyst j has the following optimal slant (without loss of generality, we assume b_1 and b_2 have zero mean and b_2 is the larger of the two):

$$s^*(d) = \gamma[(n_j - d)]/(\gamma + \beta) \dots\dots\dots(6)$$

where

$$n_j = \pm 2b_2 \mp \sqrt{\frac{4b_2^2 + \left(E\alpha - \frac{\beta\gamma}{\beta + \gamma}\sigma^2\right)}{\gamma}} \dots\dots\dots(7)$$

are the slants of the analysts. The proof of this is long and is left in the Appendix. The important result, shown also in the Appendix, is that $s_2 \geq b_2$ and $s_1 \leq b_1$. In other words, competition increases average slant:

Proposition 4a – When a stock is covered by duopolist analysts, the average slant by the analyst is higher with heterogeneous investors than with homogenous ones.

Proposition 4a – With heterogeneous investors, the average slant by duopolist analysts is higher than that by a monopolist analyst.

To put it succinctly, the key results are Propositions 4a and 4b, which say that investor heterogeneity and/or competition increases average slant. We summarize the amount of slant for the four cases in Table 1.

Table 1 – Comparing Expected slant by Analysts

This table shows the slant by analysts in a market for a particular stock covered by the analysts. Average slant by analysts for a stock is the expected value over all analyst slants. In the “homogenous investors” case, investors hold the same bias b , which is distributed $N(b, \sigma^2)$. The true signal is d . In the “heterogeneous investor” case, there are two investors, one with bias b_1 and the other b_2 . Both types of investors have utilities that put a negative weight on slant by a factor β . Heterogeneous investors’ utilities also put a negative weight on the distance between their bias and what analysts report, by a factor γ .

	Monopolist analyst	Duopolist analysts
Homogenous investors	$\gamma(b - d)/(\gamma + \beta)$	$\gamma(b - d)/(\gamma + \beta)$
Heterogeneous investors	$\gamma[(b_1 + b_2)/2 - d]/(\gamma + \beta)$	$\gamma[(n_j - d)]/(\gamma + \beta)$ where $n_j = \pm 2b_2 \mp \sqrt{\frac{4b_2^2 + \left(E\alpha - \frac{\beta\gamma}{\beta + \gamma}\sigma^2\right)}{\gamma}}$

In the next section, we test these propositions empirically

III. Econometric Model and Sample Construction

When testing if more heterogeneity among investors leads to higher average slant among analysts, a central concern is the direction of causality. To address that, we start with a simultaneous equations model and then undertake a number of robustness tests on sub-samples as well as on the time trends of the two measures of heterogeneity and average slant.

The structural equation in the simultaneous equations model relates investor heterogeneity to analyst average slant (AS) by considering other factors that might affect AS. The chief, as described in the introduction, are institutional incentives. Following Cowen, Groysberg and Healy (2003), we use the percentage of analyst institutions who are brokerages as a proxy for that:

[Structural equation]

$$AS = \beta_0 + \beta_1.Heterogeneity + \beta_2.InstitutionalIncentives + \dots [other\ exogenous\ variables] + \varepsilon$$

For the reduced form equation, we need to consider the key factor that affects investor beliefs apart from what they read from analyst reports. The primary one would be *alternative opinions* provided by in-house staff, so-called buy-side analysts. We do not have the number or quality of buy-side analysts, but we do know that mutual funds tend to have stronger and probably more opinionated buy-side analysts than regular institutional investors, since the rationale for existence of mutual funds are their proprietary ability to pick stocks². Institutional investors, on the other hand, are often funds of funds and rely heavily on asset managers and sell-side analysts. Therefore, we can construct the reduced form equation, using the percent of stock holding by mutual funds as a proxy for the strength of alternative opinions:

[Reduced form equation]

$$Heterogeneity = \beta_0 + \beta_1.AS + \beta_2.AlternativeOpinions + \dots [other\ exogenous\ variables] + \eta$$

For this paper, we limit our analysis to the analyst slants and investor beliefs about earnings per share (EPS). Apart from stock price, EPS is one of the more watched forecast variables in the industry. Unlike stock price, it does not suffer from truncation effects since it can be negative. It also does not suffer price-specific measurement problems such as those due to bid-ask bounce.

In the following sub-sections, we outline how we measure all the variables and how we construct our sample.

² This difference between mutual funds and institutional investors is also substantiated by “industry experience” of the author, who is himself a licensed securities dealer. He is also retired founder and chairman of dollarDEX, a licensed investment advisor, with offices in Hong Kong, Singapore, and Taiwan. The difference also seems to apply in the U.S. and the U.K., based on his experience as a member of the Global Advisory Council of the international London-based asset management arm of Wachovia N.A, one of the four large U.S. banks. This paper does not reflect the views of any of these establishments.

A. Measuring Average slant

Our unit of observation is a stock-time. An example observation is the EPS forecast for IBM's general stock 3 quarters before September 30, 2003 (time). "Time" is specified by both the announcement date (September 30, 2003) and the horizon (3 quarters before). For robustness, we use several measures of the average slant among analysts for an observation. The main measure follows our definition in equation (3), where we get:

$$\begin{aligned} AS &= E_d [(n - d)^2] \\ &= En^2 - 2d.En + d^2 \\ &= Vn + (En)^2 + d^2. \end{aligned}$$

Therefore, for each observation, we calculate the variance and mean of all standing forecasts (ie., including those made public earlier) for that horizon. We also get the actual EPS announced.

Following standard practice in the finance literature (e.g.: Diether, Malloy and Scherbina (2002)), another measure is to normalize our definition by taking the square root of AS and divide by the mean EPS estimate. Because the mean can be arbitrarily close to zero, we adopt the convention of limit our observations to those with mean of at least 1 cent. The robustness of this is should be evident if both measures of AS check against each other in the estimations.

The data is obtained from FirstCall, whose "fiscal period end date" (item FPE) is from 1980 through 2003. This is also the entire period available. The horizon is calculated as the difference in quarters between the estimate date (item ESTDATE) and the fiscal period end date. We use only estimates of quarterly periodicity, to control for differences in the frequency with which analysts obtain information and make forecasts.³ We start with 6,541,254 stock-time observations. Using these, we compute the AS for every observation. After dropping those without CUSIP numbers (i.e. unidentified firms) or estimates or non-quarterly forecasts, we are left with 3,326,344 observations. After accounting for various missing data for the other measures below, the final dataset has 875,851 observations, corresponding to 10,281 unique CUSIP stocks.

Table I shows the summary statistics of our AS measure, as well as other measures.

B. Measuring Investor Heterogeneity

We measure beliefs of institutional investors and mutual fund managers (collectively called "investors"), since these are the ones for whom analysts mainly write their reports.

For robustness, we use a number of measures for the heterogeneity of investor beliefs. One measure is a weighted standard deviation of percent share holdings of each investor from the mean, where the weighting is the number of shares held by the investors. Given our large population, we treat the mean as the market percent holding. Because we attach the weight to the percentage holding, this first measure reduces to just the standard deviation of the share holdings. Our second measure is to attach the weighting to the moment, rather than the raw percent. We also construct a third measure similar to the first standard deviation, but of share changes rather than share holdings. Finally, we also have a simple weighted standard deviation of share changes. While the

³ The choices of periodicity are: annual, semi-annual, and quarterly. We pick quarterly, to get more observations.

share changes are relatively comparable, the shareholdings can vary greatly with shareholdings, so we normalize these by dividing by the total shareholding held by all investors.

Table I
Summary Statistics

This table shows key variables in our dataset and their summary statistics. The data is compiled from over 6.5 million estimates from FirstCall and more than 42.8 million data items on investors' holdings from Thomson Financial. All are for analysts forecasting earnings per share (EPS) for investors, of which there are two major types: institutional and mutual funds. Each observation is a stock-time, where time is a pair of information: the "fiscal period end date" (date at which actual EPS is announced) and the forecast horizon, in quarters. "average slant" is the sum of squares of the difference between estimates and actual. Normalization is by dividing the raw AS by the mean estimate. "Investor heterogeneity" is measured primarily by the standard deviation of investors' shareholdings, or changes in shareholdings. The weighted measures are by shares. "Spread/price" is the bid-ask spread divided by stock price at the observation time. "Institutional incentives" is measured by the number of analysts who work in brokerages, which tend to be more commission-based than other types of institutions. This information is obtained by program- and hand-matching against the Hoover's Online database of brokerages. "Alternative opinions" are the sources to which investors can turn to. It is measured by the percent of shareholdings held by mutual funds rather than institutional investors, since mutual funds tend to have significantly stronger in-house buy-side analysts. "Competition" is measured by the number of analysts with standing forecasts at the observation stock-time. "Volatility" is measured by the beta. "Horizon" is the number of quarters to the fiscal period end date for an observation. "Year" are the dummies to demarcate periods of stock market booms and busts. "Market capitalization" is in thousands of dollars, at the observation stock-time.

	Number of observations	Mean	Std. Deviation	Min	Max
Fiscal period end date	875,851	3 2 1999	1,299 days	31 8 1989	31 12 2004
Average slant (AS)	875,851	7.8	1,508.9	0	802,816.0
Normalized (square root / mean)	776,238	41,172.65	5,627,685	0	3.13e+09
Normalized, limited to mean ≥ 1 ct	690,955	0.689	2.638	0	723.512
Investor heterogeneity (normalized)					
Shares SD	719,277	0.05	0.07	0	0.71
Share changes SD	710,639	0.03	0.26	0	72.92
Shares SD (weighted)	724,835	0.07	0.06	0	0.32
Share changes SD (weighted)	720,001	0.04	0.22	0	51.43
Spread / price	538,107	0.19	0.29	-15.75	17.25
Inst. incentives (% brokerages)	875,851	0.55	0.39	0	1.00
Alternative opinions (% funds)	724,835	0.19	0.19	0	1.00
Competition (number of analysts)	875,851	3	4	1	53
Volatility (beta)	413,467	0.76	0.64	-2.16	3.70
Horizon in quarters	875,851	3.13	1.91	1	20
Year					
1990-94	136,611	1	0	1	1
1995-99	364,291	1	0	1	1
2000-03	374,870	1	0	1	1
Market capitalization	715,326	3,403	16,955	0	1,047,533

We believe the measure using the sum of squares of share holdings is appropriate because investors who have a high holding of a stock are likely to believe in that stock, both when they bought it and probably even at the time of our observation, if we assume that investors rationalize their past purchases. Our measure is robust to three possible concerns. First, it might be argued that there are investors who might have negative

feelings about a stock, but cannot hold short positions because of short-sale constraints on the investment mandates of many institutional investors. This causes our measure to be downward biased. Fortunately, such constraints need not factor into our consideration because analysts do not write reports for investors with no holdings anyway – i.e., they are not even in our population, in practice and by definition.

Another concern is that it might be the change in share holding, and not the level, that better reflects investors' beliefs. As described earlier, we do have an alternate measure of the change in share holding. However, the level of holding is also justified because it is probably more often that an investor with good (bad) feelings about a stock might do nothing about her holdings of it given that she has already purchased (sold down) the stock, or because the change in feeling is not sufficiently large to justify the expected change in return in view of transaction costs. The investor might also not sell (buy) a stock even when she thinks poorly (well) of it, if she believes she can get it even cheaper (dearer) later.

The third concern is that our notion of market holding is the mean *at the time of the observation*, rather than the market holding at the time after the actual EPS is announced. This is a discrepancy from how AS is calculated. It may also confound heterogeneity of beliefs with the simple heterogeneity of sizes of investor holdings. This is a legitimate concern and we plan to incorporate this modification later.⁴ In this paper, we simply note the possibility of this confounding effect. To the extent that the market as a whole has an accurate forecast of what it should be holding, our measure will not be biased.

We obtain the share holdings from Thomson Financial's "CDA/Spectrum Institutional 13(f)" database for institutional investors and their "CDA/Spectrum Mutual Funds Holding US & Canadian Securities" database for mutual funds. We again take the entire period available, from 1980 through 2003. This consists of 42,866,417 observations, out of which 26,887,719 are institutional investors and 15,978,698 mutual funds. For each quarter-end (as denoted by item RDATE in the Thomson database), we calculate the measure of heterogeneity, as well as other measures that we describe below (such as market capitalization).

B. Measuring Institutional Incentives

In our structural equation, we need instruments *that do not explain heterogeneity in the reduced equation* (exclusion restriction). At the same time, the instruments need to be correlated with AS (identifying condition). The traditional view on bias suggests just such an instrument: distorted incentives by analysts' firm. Following Cowen, Groyberg and Healy (2003), we use the percentage of analysts who come from brokerage firms as a proxy for this measure. Brokerage firms are commonly fee-based, while other firms such as investment banks or pure research houses might have other sources of income.

The FirstCall database of estimates provides the 398 names of unique firms who make the estimates. For each stock, each time (i.e., fiscal period end date and horizon in quarters), we count the number of firms who are brokerages divided by the total number of firms who have standing estimates for that fiscal period end date. To determine whether a firm in FirstCall is a brokerage, we match each name against the 309 firms listed as "Securities Brokers & Trades" in Hoover's Online. The matching is done first by a program, using the first proper name in the FirstCall name against the Hoover's list.

⁴ We have not done it in this paper because it turns out to be enormously complicated for the datasets of the size we are dealing with here. Please see later discussion on sizes of datasets.

For example, “Lehman” in “Lehman Brothers” is used. If there is a match, the match is manually checked to see if it is spurious. Then the second proper name is used, such as “Witter” in “Dean Witter.” Third, the remaining are manually checked, using clues such as whether there is “Securities” in the name, which is almost certainly a word associated with brokerages. Of the FirstCall firms, 120 are found using our procedure.

D. Measuring Alternative Opinions

Likewise, “alternative opinions” in the reduced form equation is assumed to influence investor heterogeneity but not average slant – i.e., satisfy the exclusion restriction. Again, it is to be correlated with heterogeneity (identifying condition). We measure the presence and strength of alternative opinions by the share holding of the stock that is held by mutual funds rather than institutional investors.

The economic reasoning is that investors use alternative opinions in addition to analysts’ (average slanted) reports to form their (heterogeneous) beliefs. However we are hard put to find why these alternative opinions might systematically influence slant by analysts. While it is possible to surmise that analysts who face investors with strong buy-side analysts might be less inclined to slant, the speculation could be in the other direction too, with sell-side analysts slant more to distinguish themselves. In practice, however, it is more often that sell-side analysts, in making the same EPS forecast for more than one investor, are likely to not be overly influenced by this factor.

The data for this measure is obtained from the Thomson Financial databases mentioned earlier, where we calculate the holdings in the mutual funds database as a percentage of holdings in both the mutual funds and institutional investors databases.

E. Other Variables

Other explanations for AS and heterogeneity include: (1) the number of analysts making for each stock-time observation, (2) the volatility of the stock, (3) the horizon, (4) year effect, and (5) market capitalization at the time of the observation.

The number of analysts is a key variable in testing our propositions about industry structure. The data is obtained by counting the number from our FirstCall database, by stock and by fiscal period-horizon.

All things being equal, the more volatile a stock, the higher the average slant by analysts. For an analyst who wants to slant, a more volatile stock provides a good justification both to herself as well as to others if she is proven incorrect. We measure volatility using betas from the CRSP. Specifically, we use the “Indices” section, downloading the “Year-end Beta Deciles Assignments.” These betas are for years rather than quarters, and in matching them up to our dataset, we assume that different quarters in the same year have the same beta. This is justified since betas are supposed to be reasonably stable in the short-term.

The horizon clearly affects the AS, since by definition, AS shrinks to null by the fiscal period end date. The horizon, as noted before, is in quarters.

We also factor in a year effect, which we measure using dummies for five-year periods from 1980-84 through 1995-99. The last dummy is 2000-03. We use the year effect to account for stock market booms and busts. For example, in a boom it is easier for many analysts to slant positively, increasing AS, even if all other things are equal.

The market capitalization of a stock measures more than volatility or liquidity. We reason that the larger the market cap, the higher the likelihood that an investor will hold

the stock, regardless of belief. Therefore, the less is AS influenced by heterogeneous beliefs. The market cap is calculated for each observation by multiplying the Thomson Financial PRC (“end of quarter share price”) and SHROUT1 (“end of quarter shares outstanding in millions”) items.

F. Omitted Variables and Asymptotic Properties

There could be other factors which determine AS. For example, the skill and experience of analysts, and those of the investors, might affect accuracy of EPS forecasting. Gearing also affects EPS. Because we are looking at AS, we need to consider the variance and not the levels of these factors. There appears to be no straight-forward relationship between these and our explanatory variables, so we leave these in the disturbance term. For robustness, we also test one of these, gearing, for estimation bias using the Hausman test, and the result (unreported) is that it does not affect the bias of the parameters at the 1% significance level.

Finally, we note that given our large sample size, we can comfortably rely on the less stringent conditions needed for asymptotic consistency and efficiency, rather than the full Gauss-Markov requirements.

IV. Empirical Results

This section begins with the main results of our estimation. This is followed by further estimations to ensure robustness and to provide deeper insight into our key variables of average slant and investor heterogeneity. We first test “space,” then we test “time.” By “space,” we mean diving into sub-samples of our dataset by type of client to see how analysts cater to these different types. By “time,” we mean testing AS and investor heterogeneity using time series methods, to understand the nature of these series individually (e.g., trending, serial correlation, unit roots) and together (i.e., cointegration).

A. Main Result

We first run the simultaneous equations model using two-stage-least-squares (2SLS). Since 2SLS is less efficient and has larger standard errors than ordinary least squares (OLS), we check if there is endogeneity in the model before proceeding with it. A Hausman test results in a χ^2 of 6.22 and a p -value of 0.0126. Therefore, endogeneity is a problem at the 5% level. We should test the reduced form for the identifying condition. The coefficient on the instrument *%Funds* is -0.103 with a standard error of 0.001, and a p -value of 0.00. Likewise, the identifying condition test on the structural equation gives the instrument *%Brokerage* a coefficient of -0.409 and a standard error of 0.135. The p -value is 0.002. In both cases, the correlations are significant, so the identifying conditions are satisfied.

Table II shows the results of the 2SLS regressions. For robustness, we use the 2 different measures of average slant (AS) and 4 different measures of heterogeneity described in the earlier section. The AS measures are shown as Panels A and B, and the regressions using the 4 measures of heterogeneity are shown in the columns.

If our catering theory holds, we expect the coefficients on investor heterogeneity and competition in the structural equation to be positive when predicting average slant (AS). We might have considered an interaction term, but given that “heterogeneity” is defined

as more than 1 investor with a different bias and “competition” is defined as an industry structure with more than 1 analyst, the homogenous and monopolist cases are trivial. To confirm this, we run regressions with interaction terms and obtain lower F ratios on them. For example, the first specification in Table II has an F statistic of 3.30, but adding an interaction term reduces it to 2.75. Other specifications have the same result.

In Table II, we see that the coefficients for heterogeneity are indeed positive in the first two columns, and in both panels. However, the coefficients from the “share change” measures of heterogeneity are zero. Without re-iterating our doubts about share changes in the earlier section, we interpret this as share changes being a poor measure of heterogeneity. This also means that while the estimates are asymptotically consistent, they may have the wrong standard errors in the rightmost two columns.

Importantly, the coefficients for competition also support the theory’s prediction. Competition does not appear to reduce to average slant.

The signs for other coefficients are also either expected or not statistically significant. For example, the coefficient for volatility is positive where they are significant: higher volatility can be interpreted as providing more room for analysts to slant. Except for an statistically insignificant result in Panel B, the coefficient for horizon is also positive, as expected.

A prominent exception is “institutional incentives,” where we expect to have a positive sign: a stock covered more by commission-based brokerages is expected to have a larger AS, but the coefficients are negative, at least in Panel A. We interpret this as a measurement error, either of the dependent variable or of our measure of “percent by brokerages.” The dependant variable AS is not normalized, so that could be a source of this. However, measurement errors in the dependent variable do not compromise the validity of ordinary least squares, except they lead to bigger standard errors. If the measurement error is with on the explanatory variable, then it might lead to biased standard errors for all our explanatory variables. For example, it is admittedly counter-intuitive that analysts in brokerages are more biased than those in investment banks who underwrite stocks that the analysts cover, no matter how strong the Chinese wall is between the analysts and investment banking. Thus, a better measure might have included not just brokerages but also underwriting investment banks. For our purposes, and to ensure our conclusions hold, we run specification (5) in the rightmost column of Panel A, without “institutional incentives.” It turns out that the regression has a better F after all, and the standard errors for other explanatory variables are not compromised. Therefore, our estimation holds.

Finally, we note that the R^2 are negligible in these regressions, because of the distorting effect of the very large number of observations.

Table II
Basic Results

This table shows key results using two different measures of average slant (AS), the dependant variable, which is compiled from over 6.5 million analyst estimates of earnings per share (EPS) from FirstCall. These are shown as Panels A and B. In Panel A, AS is defined as $E(n-d)^2$ where n is an estimate and d is the actual EPS. The expectation is taken over all analysts for an observation, which is defined by a stock, a fiscal period end date (i.e., the date the actual EPS would be announced), and the forecast horizon in quarters. In Panel B, we normalize the AS calculated in Panel A by taking its square root and dividing it by the mean over all analysts for the observation. However, we limit the observations to those where the mean EPS is at least 1 cent. The four columns (1) through (4) correspond to four measures of investor heterogeneity, the key explanatory variable. These are compiled from more than 42.8 million data items on investors' holdings from Thomson Financial. "Shares SD" is the standard deviation of holdings among investors for each observation. "All shares" is the holdings among all investors for that stock at that time specified by the observation. "Weighted" is another measure where the weight is on the moment rather than the raw shareholding percent. "Institutional incentives" is measured by the percentage of analysts firms who are brokerages. All regressions are done using a simultaneous equations model with two-stage least-squares. The percentage of shareholdings held by mutual funds is used as an instrument for investor heterogeneity. Column (5) in Panel A is to test a specification without "institutional incentives," which is suspected to be mis-measured. The column shows that the regression F statistic is better than that in (1), and the significance of the other variables are uncompromised. The figures in brackets are heteroskedastic-robust standard errors.

	(1)	(2)	(3)	(4)	(5)
Panel A: Dependant variable is average slant (AS)					
Investor heterogeneity					
Shares SD / All shares	17.54 (7.38)				17.33 (4.85)
Shares SD / All shares (weighted)		11.96 (3.33)			
Share change SD			-0.000004 (0.000001)		
Share change SD (weighted)				-0.000003 (0.000001)	
Inst. incentives (% brokerages)	-0.21 (0.10)	-0.32 (0.12)	-0.37 (0.13)	-0.32 (0.12)	
Competition (number of analysts)	0.20 (0.08)	0.16 (0.07)	0.25 (0.09)	0.36 (0.12)	0.19 (0.08)
Volatility (beta)	-0.07 (0.22)	-0.21 (0.24)	0.17 (0.22)	0.38 (0.22)	-0.07 (0.22)
Horizon in quarters	0.10 (0.04)	0.06 (0.03)	0.10 (0.04)	0.16 (0.06)	0.09 (0.04)
Year					
1990-94	1.77 (0.45)	1.78 (0.45)	1.96 (0.48)	1.83 (0.45)	1.78 (0.45)
1995-99	-0.86 (0.38)	-0.71 (0.35)	-0.03 (0.23)	-0.27 (0.26)	-0.85 (0.38)
2000-03	-0.85 (0.44)	-0.81 (0.43)	0.35 (0.28)	0.43 (0.28)	-0.83 (0.43)
Market capitalization	-0.007 (0.003)	-0.006 (0.003)	0.02 (0.01)	0.06 (0.02)	-0.01 (0.00)
Intercept	-0.76 (0.26)	-0.46 (0.22)	-0.02 (0.18)	-0.51 (0.21)	-0.87 (0.29)
Number of observations	355577	356564	351337	354978	355577
F	3.31	3.33	3.15	2.86	3.60
p -value	.0005	.0005	.0008	.0022	0.0003
R^2	-	-	-	-	-

	(1)	(2)	(3)	(4)	(5)
Panel B: Dependant variable is square root of average slant divided by mean					
Investor heterogeneity					
Shares SD / All shares	3.79 (0.51)				3.86 (0.51)
Shares SD / All shares (weighted)		2.48 (0.33)			
Share change SD			-0.0000008 (0.0000001)		
Share change SD (weighted)				0.0000005 (0.0000001)	
Inst. incentives (% brokerages)	0.06 (0.01)	0.03 (0.01)	0.02 (0.01)	0.03 (0.01)	
Competition (number of analysts)	0.01 (0.00)	0.007 (0.0017)	0.022 (0.004)	0.045 (0.01)	0.014 (0.002)
Volatility (beta)	0.39 (0.01)	0.36 (0.01)	0.44 (0.02)	0.48 (0.02)	0.39 (0.01)
Horizon in quarters	0.01 (0.00)	-0.0014 (0.0028)	0.01 (0.003)	0.02 (0.005)	0.01 (0.003)
Year					
1990-94	0.50 (0.12)	0.50 (0.11)	0.56 (0.11)	0.54 (0.12)	0.49 (0.12)
1995-99	0.72 (0.12)	0.74 (0.11)	0.90 (0.11)	0.86 (0.12)	0.71 (0.12)
2000-03	0.94 (0.12)	0.94 (0.11)	1.20 (0.12)	1.22 (0.12)	0.94 (0.12)
Market capitalization	-0.0032 (0.00019)	-0.0031 (0.00019)	0.0013 (0.0008)	0.01 (0.002)	-0.003 (0.0002)
Intercept	-0.48 (0.12)	-0.41 (0.11)	-0.35 (0.12)	-0.44 (0.13)	-0.45 (0.12)
Number of observations	303,261	304,021	299,984	302,800	303,261
<i>F</i>	189	190	152	146	210
<i>p-value</i>	.0000	.0000	.0000	.0000	.0000
<i>R</i> ²	.0094	.0095	-	-	.0093

We can now address the propositions posed earlier. As expected, Proposition 1 (“With rational investors, the average slant (AS) is zero, for both monopoly and duopoly industrial structures”) is not supported, since the coefficient for “competition” is significantly non-zero. Proposition 2 (“With behavioral homogenous investors, the average slant (AS) the same for both monopoly and duopoly industrial structures.”) is weakly supported because estimations with interactions between competition and heterogeneity have very weak statistical significance, even though the signs are as expected. Proposition 3 (“When a stock is covered by a monopolist analyst, the average slant by the analyst is a function of the mean bias among the investors, whether for the homogenous or heterogeneous cases.”) is supported by construction. The main tests, for Proposition 4a (“When a stock is covered by duopolist analysts, the average slant by the analyst is higher with heterogeneous investors than with homogenous ones.”) and Proposition 4a (“With heterogeneous investors, the average slant by duopolist analysts is higher than that by a monopolist analyst.”) are supported, the former with positive coefficients for “heterogeneity” and the latter with the same for “competition.”

B. Sub-samples by Investor Type

We re-do our estimation using sub-samples corresponding to different types of investors. This provides some sense of the robustness of our results down to investor-type level, and offers insight into how analysts might cater to investors differently. The results are shown in Table III, with Panel A showing the different types of institutional investors and Panel B the types of mutual funds.

Investor types are from the same Thomson Financial datasets for our main estimation. The observations are now confined to the specific types. For example, for “banks,” the AS is calculated for stocks held only by banks, and investor heterogeneity is calculated for only share holdings by banks.

In Panel A, we see that the main results from the previous sub-section broadly hold. We expect institutional investors with strong teams of in-house analysts, such as banks and investment managers, to buck our theory more, since analysts are less likely to cater to clients with already strong analyses. This holds up in the result. Heterogeneity is even negative for investment managers who, like mutual fund managers, are probably the ones with the strongest in-house teams. This means that sell-side analysts have no incentive to cater to them, as the formers’ analyses are likely to be discounted or even subject to a “not invented here” view. For banks and investment managers, the other main variable “competition” also has negative coefficients. With these investor types, competition among analysts is less likely to result in them separating out to cater to the largest, and more eager to be accurate in the face of an expert audience. Other variables all have the right sign or are otherwise statistically insignificant.

In Panel B, we expect that mutual funds in general are experts and in-house buy-side analysts, as with the banks and investment managers in Panel A. We therefore expect the size of the coefficients for heterogeneity and competition to be either closer to zero than those in Panel A, or downright negative. This is broadly supported. It also conforms with intuition that the more aggressive funds such as “aggressive growth” or “growth, income” have more catering, with their bigger coefficients on heterogeneity. There are more views in these areas, than say in the staid bond market, and sell-side analysts can afford to cater or have a role to play by confirming buy-side analysts’ work.

C. Time Series Analysis

In this section, we characterize AS and heterogeneity as time series, and see how they might be related. This provides another lens to view the relationship between the two, serving both as a robustness test and a means to gain deeper insight.

Here is the roadmap. We begin with some basic specifications test, to check if the two are related, and if so, whether the relationship is spurious because they have time trends. We also check for serial correlation, and correct for these. To be sure, we also check these series for unit roots, since that might invalidate our estimations. There is some evidence of that, so we next look at whether they might have a cointegration factor, even if they are individually integrated of higher orders, $I(n)$. The result shows that they are indeed cointegrated. This leads us to explore an even tighter relationship: whether one series Granger-causes the other.

Table III
Results by Type of Investor

This table shows results by the type of investor. Panel A shows institutional investors, and Panel B shows mutual funds. The dependant variable is average slant (AS) in both cases. AS is defined as $E(n-d)^2$ where n is an estimate and d is the actual EPS. The expectation is taken over all analysts for an observation, which is defined by a stock, a fiscal period end date (i.e., the date the actual EPS would be announced), and the forecast horizon in quarters. The information for AS is compiled from over 6.5 million analyst estimates of earnings per share (EPS) from FirstCall. Investor "heterogeneity," the key explanatory variable, is compiled from more than 42.8 million data items on investors' holdings from Thomson Financial. We measure it using "Shares SD," the standard deviation of holdings among investors for each observation, which is also the weighted percent holding from the market holding. "All shares" is the holdings among all investors for that stock at that time specified by the observation. "Institutional incentives" is measured by the percentage of analysts firms who are brokerages. All regressions are done using a simultaneous equations model with two-stage least-squares. The percentage of shareholdings held by mutual funds is used as an instrument for investor heterogeneity. The figures in brackets are heteroskedastic-robust standard errors. "Dropped" in Panel B means the variable is dropped from the regression because of multicollinearity.

Panel A: Institutional Investors				
	Banks	Insurance companies	Investment managers	Independent advisors
Heterogeneity (Shares SD/All shares)	11.85 (10.67)	221.38 (108.09)	-1587.99 (1456.34)	11.92 (6.11)
Inst. incentives (% brokerages)	-0.40 (0.35)	-1.09 (0.59)	-2.57 (2.31)	-0.56 (0.31)
Competition (number of analysts)	-0.02 (0.01)	0.37 (0.24)	-1.68 (1.75)	0.02 (0.01)
Volatility (beta)	0.41 (0.24)	-0.40 (0.80)	-8.95 (7.68)	0.44 (0.19)
Horizon in quarters	-0.02 (0.02)	0.18 (0.16)	-1.68 (1.65)	0.04 (0.03)
Year				
1990-94	0.97 (0.51)	-0.10 (1.12)	34.87 (31.50)	1.29 (0.66)
1995-99	0.12 (0.14)	-4.33 (2.56)	60.16 (56.30)	-0.04 (0.33)
2000-03	0.43 (0.19)	-1.93 (1.78)	-1.24 (8.37)	0.37 (0.34)
Market capitalization	0.00 (0.00)	-0.02 (0.01)	-0.06 (0.05)	0.00 (0.00)
Intercept	-0.36 (0.33)	-1.17 (1.07)	40.53 (37.74)	-0.86 (0.47)
Number of observations	86,072	57,899	48,812	91,619
F	1.17	0.72	0.33	1.81
p-value	.3090	.6915	.9648	.0602
R ²	-	-	-	-

Panel B: Mutual Funds							
	Interna- tional	Aggressive growth	Growth	Growth, income	Bond & preferred	Balanced	Metals
Heterogeneity	1.90	341.60	14.01	115.50	13.32	-5.93	-0.15
(SharesSD/Allshares)	(2.31)	(360.84)	(7.92)	(103.57)	(14.26)	(7.64)	(0.89)
Inst. incentives (% brokerages)	0.03 (0.05)	0.19 (0.92)	-0.07 (0.18)	0.10 (0.33)	-2.18 (1.71)	0.01 (0.03)	0.03 (0.11)
Competition (no. of analysts)	0.01 (0.01)	1.73 (1.89)	0.50 (0.30)	0.41 (0.40)	-0.03 (0.03)	0.00 (0.00)	0.00 (0.01)
Volatility (beta)	0.04 (0.04)	5.46 (5.81)	-0.49 (0.50)	0.37 (0.63)	1.50 (1.35)	-0.01 (0.03)	-0.02 (0.02)
Horizon in quarters	0.02 (0.01)	1.28 (1.43)	0.19 (0.13)	0.23 (0.19)	-0.12 (0.14)	0.01 (0.01)	0.01 (0.00)
Year							
1990-94	(dropped) (0.00)	8.93 (29.09)	4.34 (2.35)	13.72 (14.30)	(dropped) (0.00)	(dropped) (0.00)	(dropped) (0.00)
1995-99	0.04 (0.03)	2.97 (27.78)	0.55 (1.24)	11.16 (12.46)	-0.30 (0.45)	-0.28 (0.28)	-0.17 (0.09)
2000-03	0.06 (0.03)	20.64 (34.93)	0.62 (1.25)	13.79 (14.35)	1.55 (1.24)	-0.36 (0.34)	-0.20 (0.14)
Market capitalization	0.00 (0.00)	0.11 (0.11)	-0.02 (0.01)	-0.01 (0.01)	0.00 (0.00)	0.00 (0.00)	0.00 (0.01)
Intercept	-0.18 (0.23)	-61.94 (71.81)	-4.14 (2.53)	-20.19 (19.27)	-0.23 (0.89)	0.45 (0.39)	0.19 (0.22)
Number of obs.	26,214	45,506	94,987	19,375	451	22,707	223
<i>F</i>	4.82	0.29	1.23	0.35	0.87	5.30	6.92
<i>p-value</i>	.0000	.9775	.2702	.9575	.6219	.0000	.0000
<i>R</i> ²	-	-	-	-	.0205	-	.3738

Table IV shows our first time series specifications for a simple. The time series is constructed from the median AS and heterogeneity over all observations for each fiscal period end date (i.e., date on which actual EPS is announced). We test if a time series relationship between AS and heterogeneity is spurious. A comparison of the “no trend” and “time trend” columns shows that the statistical significance of our key variables, heterogeneity and competition, change very little, if at all, with the time trend. We conclude that time trend is not a cause for spurious correlation, if any.

Panel B shows tests for serial correlation. The left two columns show the results of an AR(1) test, where we first regress (using 2SLS and the “percent mutual funds” instrument) to obtain the predicted residuals. We then regress the predicted residuals on its lag, as well as the explanatory variables. The significance of the coefficient on the lagged residual suggests there is serial correlation. Indeed, the serial correlation might have more than 1 lag, as the middle two columns show. The *F* and Breusch-Godfrey *LM* statistic are very significant. We therefore correct for serial correlation in the rightmost two columns, using Cochrane-Orcutt FGLS. The result continues to support the propositions of the theory.

Despite the lack of time as a possibility for spurious correlation, given the serial correlation, we should also test if the series themselves could have unit roots. Table V shows the results. The model for the AS test is:

$$\Delta AS_t = \beta_0 + \beta_1 AS_{t-1} + \varepsilon_t$$

Table IV
Times Series

This table shows the time series relationships between average slant and investor heterogeneity. Panel A shows tests for time trends, and Panel B shows tests for serial correlation. The dependant variable is average slant (AS) in both cases. AS is defined as $E(n-d)^2$ where n is an estimate and d is the actual EPS. The expectation is taken over all analysts for an observation, which is defined by a stock, a fiscal period end date (i.e., the date the actual EPS would be announced), and the forecast horizon in quarters. The information for AS is compiled from over 6.5 million analyst estimates of earnings per share (EPS) from FirstCall. Investor "heterogeneity," the key explanatory variable, is compiled from more than 42.8 million data items on investors' holdings from Thomson Financial. We measure it using "Shares SD," the standard deviation of holdings among investors for each observation, which is also the weighted percent shareholding from the market shareholding. "Weighted" is another measure where the weight is on the moment rather than the raw shareholding percent. "All shares" is the holdings among all investors for that stock at that time specified by the observation. The time series is constructed from the median AS and heterogeneity over all observations for each fiscal period end date (i.e., date on which actual EPS is announced). Panel B tests for serial correlation without assuming strict exogeneity. We assume only contemporaneous exogeneity between the residuals and the regressors in the same period. The residual is for a simultaneous equations model, using a 2SLS estimation with the percentage of shareholdings by mutual funds as an instrument for heterogeneity. The dependant variable in the tests shown is the residual, and is estimated using the same 2SLS with the same instrument. The figures in brackets are heteroskedastic-robust standard errors.

	Panel A – Time Trends			
	No trend		Time trend	
	SharesSD/ AllShares	Weighted	SharesSD/ AllShares	Weighted
Heterogeneity	.06	.11	.06	0.11
(Shares SD / All shares)	(.06)	(.04)	(.06)	(0.04)
Competition	.0003	.0001	.0003	.0002
(no. of analysts)	(.0005)	(.0005)	(.0006)	(.0005)
Volatility (beta)	-.0059	-.0054	-.0060	-.0056
	(.0013)	(.0013)	(.0013)	(.0013)
Horizon in quarters	.0011	.0007	.0011	.0009
	(.0004)	(.0004)	(.0004)	(.0004)
Market capitalization	-.0042	-.0031	-.0042	-.0026
	(.0018)	(.0015)	(.0019)	(.0019)
Time			-.0000004	-.0000044
			(-.000008)	(-.000008)
Intercept	.0061	.0026	.0061	.0023
	(.0026)	(.0024)	(.0026)	(.0026)
Number of obs.	185	185	185	185
<i>F</i>	25.68	25.99	21.35	21.98
<i>p-value</i>	0	0	0	0
<i>R</i> ²	.318	.337	.318	0.339

Panel B – Serial Correlation						
	AR(1), without assuming strict exogeneity		Additional lags		Cochrane-Orcutt FGLS	
	SharesSD/ AllShares	Weighted	SharesSD/ AllShares	Weighted	SharesSD/ AllShares	Weighted
Heterogeneity	1.29 (0.79)	1.68 (0.81)	1.52 (0.48)	1.70 (0.50)	0.21 (0.25)	0.19 (0.19)
Inst. incentives (% brokerages)	0.00 (0.01)	0.00 (0.02)	0.01 (0.01)	0.00 (0.01)	-0.01 (0.01)	-0.01 (0.01)
Competition (no. of analysts)	0.00 (0.00)	0.00 (0.01)	0.00 (0.00)	-0.01 (0.00)	0.00 (0.00)	0.00 (0.00)
Volatility (beta)	0.01 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)
Horizon in quarters	0.00 (0.00)	-0.01 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
Time	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
Market capitalization	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
Alternative opinion (% funds)					-0.24 (0.09)	-0.21 (0.10)
Residuals lagged 1	0.43 (0.06)	0.37 (0.07)	0.19 (0.07)	0.01 (0.07)		
Residuals lagged 2			0.07 (0.05)	0.10 (0.05)		
Residuals lagged 3			0.71 (0.06)	0.63 (0.07)		
Residuals lagged 4			-0.06 (0.07)	0.14 (0.08)		
Intercept	-0.04 (0.04)	-0.08 (0.05)	-0.06 (0.03)	-0.08 (0.04)	0.09 (0.03)	0.09 (0.03)
Number of obs.	184	184	181	181	184	184
<i>F</i>	7.41	8.57	60.55	79.84	7.82	7.82
<i>p-value</i>	.000	.000	.000	.000	.000	.000
<i>R</i> ²	.320	.419	.773	.826	.122	.123
Test on all residuals						
<i>F</i>			118.34	108.27		
<i>Breusch-Godfrey LM</i>			136.9	146.3		
<i>p-value</i>			.000	.000		

The test with 2 lags, for example, is:

$$\Delta AS_t = \beta_0 + \beta_1 AS_{t-1} + \beta_2 \Delta AS_{t-1} + \beta_3 \Delta AS_{t-2} + \varepsilon_t$$

For AS, although the “base” case rejects the possibility of a unit root at 1% significance, the test with lags fail to reject the null hypothesis of a unit root. As to the number of lags required, a test on just the fourth lag has an $F(1,174) = 0$ and a p -value of 0.8498. The joint test for the third and fourth lags gives an $F(2,174) = 4.07$ and a p -value of 0.0187. Finally, a joint test for the second through fourth legs results in $F(3, 174) = 20.66$ and a p -value of 0.0000. This suggests that three lags are needed. We have earlier shown that a time trend seems unimportant in predicting AS using other explanatory variables. We now test whether it is significant when the explanatory variables are all AS and AS lags. The result confirms that the time trend does not change our conclusion that a unit root cannot be rejected. Also, time has a significant coefficient. However, the conclusion

on the number of lags needed (3) is also unchanged. The fourth lag has a p -value of 0.7625, but a joint test of the third and fourth has a p -value of 0.0996.

The test for unit root in the heterogeneity series follows the same pattern: while a unit root is rejected with the simple model, it cannot be rejected with lags (and the time trend appears significant). In the regression with time trend, the fourth lag has an $F(1,173) = 8.54$ and p -value of 0.0039, and both the third and fourth lags have $F(2,173) = 1.02$ and p -value of 0.0152. Therefore, the fourth lag appears to be significant here.

Table V
Unit Roots

This table shows the unit root tests for average slant (AS, scaled using square root) and investor heterogeneity. The time series is constructed from the median AS and heterogeneity over all observations for each fiscal period end date (i.e., date on which actual EPS is announced). The model for the AS test is $\Delta AS_t = \beta_0 + \beta_1 \cdot AS_{t-1} + \varepsilon$. The test with 2 lags, for example, is $\Delta AS_t = \beta_0 + \beta_1 \cdot AS_{t-1} + \beta_2 \cdot \Delta AS_{t-1} + \beta_3 \cdot \Delta AS_{t-2} + \varepsilon$. The analogous model is used for test of the heterogeneity series. "D-F crit. value" is the Dickey-Fuller critical value for tests with time and no time trends. The figures in brackets are heteroskedastic-robust standard errors.

	Δ Average slant (AS, scaled with square root)			Δ Heterogeneity (SharesSD Weighted / AllShares)			
	Base	With lags	Time trend	Base	With lags	Time trend	
AS	-0.10	-0.36	-0.02	Heterogeneity	-0.53	-0.01	-0.03
Robust std. err	(0.03)	(0.27)	(0.03)	Robust std. err	(0.12)	(0.03)	(0.03)
t	-3.39	-1.36	-0.92	t	-4.58	-0.32	0.98
D-F crit. value				D-F crit. value			
1%	-3.43	-3.43	-3.96	1%	-3.43	-3.43	-3.96
5%	-2.86	-2.86	-3.41	5%	-2.86	-2.86	-3.41
10%	-2.57	-2.57	-3.12	10%	-2.57	-2.57	-3.12
Δ AS lagged 1		-0.42	-0.46	Δ hetero. lagged		-0.93	-1.00
		(0.09)	(0.09)	1		(0.08)	(0.08)
Δ AS lagged 2		-0.34	-0.40	Δ hetero.lagged 2		-0.96	-1.07
		(0.10)	(0.10)			(0.11)	(0.11)
Δ AS lagged 3		0.21	0.15	Δ hetero.lagged 3		-0.10	-0.19
		(0.09)	(0.10)			(0.11)	(0.10)
Δ AS lagged 4		0.01	-0.02	Δ hetero. lagged		-0.10	-0.15
		(0.08)	(0.08)	4		(0.05)	(0.05)
Time			-0.00003	Time			-0.0001
			(0.00002)				(.000003)
Intercept	0.01	0.00	0.005	Intercept	0.02	0.00	0.0004
	(0.00)	(0.00)	(0.002)		(0.01)	(0.00)	(0.002)
Number of obs.	184	180	182	Number of obs.	184	180	180
F	11.48	16.36	14.22	F	21.01	152.28	132.37
p -value	.0009	.0000	.0000	p -value	.0000	.0000	.0000
R^2	.0546	.3737	.3912	R^2	.3026	.8161	.8270

What the unit root tests show is that there is another source for spurious correlation. However, since two unit root series might still be related by cointegration, we test for that. We first obtain the residuals of the model:

$$AS_t = \beta_0 + \beta_1 \cdot Heterogeneity_t + u$$

The predicted residuals are then regressed on their difference, with lag:

$$u_t = \beta_0 + \beta_1 \cdot \Delta u_{t-1} + \varepsilon$$

The result is in Table VI. We tested four scenarios: with and without time trend, and with the two measures of investor heterogeneity. Happily, the two series have a significant cointegration factor, in all four scenarios. Therefore, even with a time view, AS and investor heterogeneity are related, despite having unit roots.

We next test if AS and investor heterogeneity might have another relationship: Granger causality. To do that, we use a vector autoregressive model (VAR) that consists of two equations. For example, the ones for two lags are:

$$\begin{aligned}
 AS_t &= \beta_0 + \beta_1.AS_{t-1} + \beta_2.AS_{t-2} + \\
 &\quad \beta_1.Heterogeneity_{t-1} + \beta_2.Heterogeneity_{t-2} + \varepsilon. \\
 \text{and} \quad Heterogeneity_t &= \beta_0 + \beta_1.AS_{t-1} + \beta_2.AS_{t-2} + \\
 &\quad \beta_1.Heterogeneity_{t-1} + \beta_2.Heterogeneity_{t-2} + \varepsilon.
 \end{aligned}$$

We test using various scenarios: different measures of AS and investor heterogeneity, as well as with and without conditioning the VAR on the other explanatory variables. The results are qualitatively the same, so we report only a simple one, in Table VII.

Table VI
Cointegration

This table shows the result of a cointegration test for average slant (AS, scaled using square root) with investor heterogeneity. The time series is constructed from the median AS and heterogeneity over all observations for each fiscal period end date (i.e., date on which actual EPS is announced). We first obtain the residuals of the regression $AS_t = \beta_0 + \beta_1.Heterogeneity_t + u_t$. The predicted residuals are then regressed on a lag of its difference: $u_t = \beta_0 + \beta_1.\Delta u_{t-1} + \varepsilon$. Investor "heterogeneity," the key explanatory variable, is compiled from more than 42.8 million data items on investors' holdings from Thomson Financial. We measure it using "Shares SD," the standard deviation of holdings among investors for each observation, which is also the weighted percent shareholding from the market shareholding. "Weighted" is another measure where the weight is on the moment rather than the raw shareholding percent. "All shares" is the holdings among all investors for that stock at that time specified by the observation. The figures in brackets are heteroskedastic-robust standard errors.

	Base		Time trend	
	SharesSD/ AllShares	Weighted	SharesSD/ AllShares	Weighted
Residuals lagged 1	-0.17 (0.06)	-0.27 (0.06)	-0.22 (0.06)	-0.27 (0.06)
Intercept	0.000003 (0.001068)	-0.0002 (0.0012)	-0.0001 (0.0011)	-0.0002 (0.0012)
Number of observations	184	184	185	184
F	9.41	23.39	13.65	23.35
p-value	.0025	.0000	.0003	.0000
R ²	.0864	.1317	.0.1056	.0.1315

Table VII shows that while two lags are insufficient, three lags are enough to explain Granger causality. Indeed, causality happens in both directions, as we would expect. The *p*-value for the "heterogeneity Granger-causes AS with 3 lags" is about 0.065, while that for "AS Granger-causes heterogeneity with 3 lags" is about 0.015.

Table VII
Vector Autoregressive Model and Granger Causality

This table shows the tests for Granger causality between average slant (AS, scaled by taking square root) and investor heterogeneity. The time series is constructed from the median AS and heterogeneity over all observations for each fiscal period end date (i.e., date on which actual EPS is announced). The vector autoregressive model (VAR) consists of two equations. For example, the one for AS (for two lags) is $AS_t = \beta_0 + \beta_1.AS_{t-1} + \beta_2.AS_{t-2} + \beta_3.Heterogeneity_{t-1} + \beta_4.Heterogeneity_{t-2} + \varepsilon$. The one for heterogeneity (again for two lags) is $Heterogeneity_t = \beta_0 + \beta_1.AS_{t-1} + \beta_2.AS_{t-2} + \beta_3.Heterogeneity_{t-1} + \beta_4.Heterogeneity_{t-2} + \varepsilon$. The figures in brackets are heteroskedastic-robust standard errors.

Dependant variable:	AS		Heterogeneity	
	2 lags	3 lags	2 lags	3 lags
Heterogeneity lagged 1	-.628 (.931)	-.692 (.801)	1.298 (.303)	.119 (.098)
Heterogeneity lagged 2	1.093 (1.071)	1.113 (.915)	-.349 (.367)	.219 (.129)
Heterogeneity lagged 3		.473 (.912)		.770 (.119)
AS lagged 1	.572 (.078)	.393 (.082)	-.014 (.030)	.008 (.013)
AS lagged 2	.316 (.088)	.022 (.092)	.015 (.033)	-.033 (.013)
AS lagged 3		.482 (.080)		.014 (.012)
Intercept	-.004 (.006)	-.014 (.005)	.001 (.002)	-.002 (.001)
Number of observations	183	182	183	182
<i>F</i>	252.13	222.96	48.95	266.67
<i>p-value</i>	0	0	0	0
<i>R</i> ²	.7941	.8278	.1105	0.8976
<i>F</i> test on casual variables				
Heterogeneity	0.97	0.46		
<i>p</i> value	.3797	.0648		
AS			0.12	0.56
<i>p</i> value			.8853	.0154

The summary of these empirical results is that various robustness tests show that heterogeneity does lead to higher average slant (AS).

V. Conclusion

We begin this paper by positing the opposite of how conventional wisdom thinks about analyst bias, that it is the result of distorted incentives by “the system” – upstream factors like the analysts’ employers and indirectly, the companies they cover. We suggest that analysts are influenced by downstream investors, the purported victims of analyst bias. There is already a theory of media bias that has this flavor, where newspapers cater to their audiences. We adapt this theory, and use an enormous database of over 50 million elements in aggregate to compute average slant and investor heterogeneity. The empirics support a catering theory of analyst bias.

We can think of a number of ways in which this research can be extended. A particularly obvious avenue is to explore the extent to which investors are really influenced by analyst bias. This is equivalent to estimating the reduced form equation in our simultaneous equations model. Given the results here, our hypothesis is that

investors are influenced significantly less than conventionally thought. If true, this would make some of the previous research on analyst bias moot.

Another avenue of research is to see what this new dynamic between investor and analyst forecast mean for stock market efficiency. For example, it is possible that more heterogeneity among investors leads to bigger average slant, which in turn makes the price discovery harder. On the other hand, if investors are not fooled by average slant especially if they are part of the cause of it anyway, then the balance tips the other way. It is fruitful to develop theories and undertake empirics to see which effect dominates.

Appendix

Lemma 1 – An investor reading a report with slant $s = \gamma[(n - d)/(\gamma + \beta)]$ has expected utility $EU = E\alpha - \beta\gamma/(\gamma + \beta)[\sigma^2 + b^2] - \gamma^2/(\gamma + \beta)[n - b]^2$.

Proof. The expected utility is:

$$E\alpha - \beta \int_d \left(\frac{\gamma}{\beta + \gamma} (n - d) \right)^2 - \beta \int_d \left(d + \frac{\gamma}{\beta + \gamma} (n - d) - b \right)^2.$$

The first integral is $-\beta[\gamma/(\gamma + \beta)]^2(n^2 + \sigma^2)$ since $Ed = 0$ and $Ed^2 = \sigma^2$. The second integral is, likewise:

$$-\gamma \left(\frac{\beta}{\beta + \gamma} \right)^2 \sigma^2 + \left(\frac{\gamma}{\beta + \gamma} \right)^2 n^2 + b^2 - 2 \frac{\gamma}{\beta + \gamma} nb.$$

Therefore, the expected utility is:

$$\begin{aligned} E\alpha - \frac{\beta\gamma}{\beta + \gamma} \sigma^2 - \lambda b^2 - 2 \frac{\gamma^2}{\beta + \gamma^2} nb \\ &= E\alpha - \frac{\beta\gamma}{\beta + \gamma} (\sigma^2 + b^2) - \lambda b^2 + \frac{\beta\gamma}{\beta + \gamma} b^2 - \frac{\gamma^2}{\beta + \gamma^2} n^2 + 2 \frac{\gamma^2}{\beta + \gamma} nb \\ &= E\alpha - \frac{\beta\gamma}{\beta + \gamma} (\sigma^2 + b^2) - \frac{\gamma^2}{\beta + \gamma^2} b^2 - \frac{\gamma^2}{\beta + \gamma^2} n^2 + 2 \frac{\gamma^2}{\beta + \gamma} nb \\ &= E\alpha - \frac{\beta\gamma}{\beta + \gamma} (\sigma^2 + b^2) - \frac{\gamma^2}{\beta + \gamma^2} (b^2 + n^2 - 2nb). \blacksquare \end{aligned}$$

Equations (6) and (7) - Behavioral heterogeneous investors.

Proof. Recall that this is the analysis for behavioral heterogeneous investors with duopolist analysts. Also recall the set-up: without loss of generality, we assume b_1 and b_2 have zero mean and b_2 is the larger of the two. We assert that:

$$s^*(d) = \gamma[(n_j - d_j)/(\gamma + \beta)] \dots\dots\dots (6)$$

where

$$n_j = \pm 2b_2 \mp \sqrt{\frac{4b_2^2 + \left(E\alpha - \frac{\beta\gamma}{\beta + \gamma} \sigma^2 \right)}{\gamma}} \dots\dots\dots (7)$$

We prove this by backward induction, starting with the “opportunity cost” of an investor reading 2 analyst reports with positions n_1 and n_2 . Then we equate this cost with the utility of the indifferent investor, the one whose bias is at the zero mean.

Using Lemma 1, an investor with bias b_j and an analyst report with slant from equation (6) has expected utility:

$$E\alpha - \frac{\beta\gamma}{\beta + \gamma} (\sigma^2 + b_j^2) - \frac{\gamma^2}{\beta + \gamma^2} (n_i - b_j)^2 .$$

Therefore, the opportunity cost of reading one report versus another is simply the difference, denoted Δ :

$$\begin{aligned} \Delta &= \gamma(n_1 - b_j)^2 - \gamma(n_2 - b_j)^2 \\ &= \gamma(n_1 - n_2)(n_1 + n_2 - 2b_j) \dots\dots\dots (8) \end{aligned}$$

Our choice of n_1 and n_2 in the proposed equilibrium guarantees that this satisfies the participation constraint of the marginal investor, whose bias b is zero. The postulated strategies are symmetric, $n_1 = -n_2$, so from equation (8), Δ is just $2b_j\gamma(n_2 - n_1)$.

Now we equate Δ with the utility of the marginal investor, whose b is zero:

$$2b_j\gamma(n_2 - n_1) = E\alpha - \frac{\beta\gamma}{\beta + \gamma} \sigma^2 - \frac{\gamma^2}{\beta + \gamma^2} n_i^2 .$$

Noting again that $n_1 = -n_2$ and using $j = 2$ for the left-hand side, and using $i = 2$ for the right-hand side, we get:

$$n_2 = \pm 2b_2 \mp \sqrt{\frac{4b_2^2 + \left(E\alpha - \frac{\beta\gamma}{\beta + \gamma} \sigma^2 \right)}{\gamma}}$$

The analogous result holds for n_1 . ■

References

Baker, Malcolm, and Jeffrey Wurgler, 2002, Market Timing and Capital Structure, *Journal of Finance* 57, 1-32.

Baker, Malcolm, and Jeffrey Wurgler, 2003, A Catering Theory of Dividends, *National Bureau of Economic Research, Inc, NBER Working Papers: 9542*.

Brown, Lawrence D., and Michael S. Rozeff, 1978, The Superiority of Analyst Forecasts as Measures of Expectations: Evidence from Earnings, *Journal of Finance* 33, 1-16.

Cowen, Amanda, Boris Groysberg, and Paul M. Healy, 2003, Which Types of Analyst Firms Make More Optimistic Forecasts?, Harvard Business School Working Paper (Harvard Business School, Boston).

Diether, Karl B., Christopher J. Malloy, and Anna Scherbina, 2002, Differences of Opinion and the Cross Section of Stock Returns, *Journal of Finance* 57, 2113-41.

Festinger, Leon, 1957, A Theory of Cognitive Dissonance, *Stanford University Press*.

- Graham, John R., 1999, Herding among Investment Newsletters: Theory and Evidence, *Journal of Finance* 54, 237-68.
- 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.
- Hong, Harrison, and Jeffrey D. Kubik, 2003, Analyzing the Analysts: Career Concerns and Biased Earnings Forecasts, *Journal of Finance* 58, 313-51.
- Kahneman, Daniel, 2003, A Psychological Perspective on Economics, *American Economic Review* 93, 162-68.
- Mastrapasqua, Frank, and Steven Bolten, 1973, A Note on Financial Analyst Evaluation, *Journal of Finance* 28, 707-12.
- Mullainathan, Sendhil, and Andrei Shleifer, 2003, The Market for News, MIT Working Paper (MIT, Cambridge).
- O'Brien, Patricia C., 2003, Discussion--The Predictive Value of Analyst Characteristics, *Journal of Accounting, Auditing & Finance* 18, 649-53.
- Shleifer, Andrei, and Robert W. Vishny, 2003, Stock Market Driven Acquisitions, *Journal of Financial Economics* 70, 295-311.