The Pretty Good Analyst *

Dan Bernhardt
University of Illinois
University of Warwick

Chi Wan
University of Massachusetts, Boston

Zhijie Xiao
Boston College

July 10, 2014

Abstract

We develop and estimate a dynamic model of the recommendation decisions made by financial analysts. We provide overwhelming evidence that analysts revise recommendations reluctantly, introducing frictions to avoid frequent revisions. Publicly-available data matter far less for explaining recommendation dynamics than do the recommendation frictions and the long-lived information that analysts acquire that the econometrician does not observe. Estimates suggest that analysts structure recommendations strategically to generate profitable order flow from retail traders. We provide extensive evidence that our model describes how investors believe analysts make recommendations, and that investors value private information revealed by analysts' recommendations.

JEL classification:
Keywords: Financial Analyst Recommendations; Recommendation Revisions; Recommendation Stickiness, Asymmetric Frictions; Duration; MCMC methods.

*We thank Murillo Campello, Roger Koenker, Jenny Tucker and seminar participants at the University of Illinois, Queen's University, and the University of Florida for helpful comments. Emails: Dan Bernhardt, danber@illinois.edu; Chi Wan, chi.wan@umb.edu; Zhijie Xiao, xiaoz@bc.edu.

1 Introduction

One of the most important services that financial analysts provide is to make recommendations to retail and institutional customers about which stocks to purchase, and which ones to sell. Brokerage houses want to employ financial analysts who provide recommendations on which investors can profit, thereby generating profitable trading activity for the brokerage house. Many researchers (e.g., Womack (1996), Barber et al. (2001), Jegadeesh et al. (2004), Ivkovic and Jegadeesh (2004)) have documented the profitability and informativeness of various measures of recommendations and recommendation changes.

In this line, one can contemplate an "idealized" financial analyst who first exhaustively gathers and evaluates information from public and private sources about a set of companies to form assessments about their values. Ignoring for the moment what enters this valuation assessment, an idealized analyst would then compare his value assessment with the stock's price, and issue buy or sell recommendations to his investor audience on that basis. Thus, an idealized analyst employing a five-tier rating system would issue "Strong Buy" recommendations for the most under-valued stocks, whose value-price differentials, $\frac{V-P}{P}$, exceeded a high critical cutoff, μ_5. The analyst would establish progressively lower cutoffs, μ_4, μ_3 and μ_2, that determine "Buy", "Hold", "Sell" and "Strong Sell" recommendations, so that, for example, the analyst would issue Buy recommendations for value-price differentials between μ_5 and μ_4, and strongly advise customers to sell stocks with the worst value-price differentials below μ_2.

The data *violently* reject the hypothesis that financial analysts form recommendations in this way. To understand why, observe that sometimes a stock's value-price differential will be close to a critical cutoff, in which case slight fluctuations in price relative to value can cause an analyst's assessment to alternate above and below the cutoff, causing the analyst to revise recommendations repeatedly (see Figure 1). In practice, analysts infrequently revise recommendations. This reluctance to revise likely reflects that customers may question the ability of an analyst who repeatedly revised recommendations.

In this paper, we develop and estimate a dynamic model of value assessment and recommendations issued by a "pretty good" financial analyst. A pretty good analyst assesses value in the same way as the idealized analyst, and when initiating coverage of a stock, the analyst makes an initial recommendation on the same basis, for example issuing a Strong Buy recommendation if and only if his assessment of the value-price differential, $\frac{V-P}{P}$, exceeds the high cutoff μ_5. However, the analyst understands how customers may interpret frequent revisions,

Figure 1: Idealized analyst

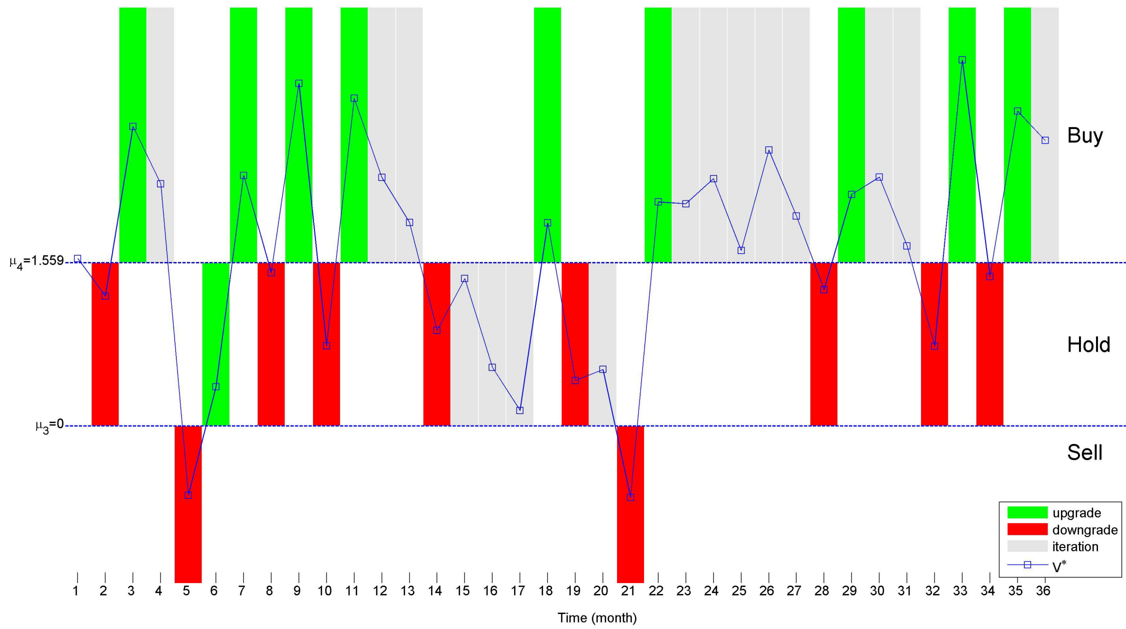

Recommendation cutoffs, and sample valuation and recommendation paths over a 36-month period for an idealized analyst employing a three-tier rating system whose private information is transient.

and only reluctantly revises recommendations: a pretty good analyst only downgrades a recommendation if the value-price differential falls far enough below the critical cutoff, and only upgrades a recommendation if the value-price differential rises far enough above the critical cutoff. Concretely, a pretty good analyst downgrades a recommendation from a Buy only if the stock's value-price differential falls below $\mu_4 - \delta_{4\downarrow}$, and upgrades a recommendation from a hold only if the differential rises above $\mu_4 + \delta_{4\uparrow}$ (see Figure 2). This model nests the "idealized" financial analyst, who sets recommendation revision frictions, $\delta_{k\uparrow}$ and $\delta_{k\downarrow}$, of zero.

We face several econometric challenges. Analysis is complicated by the fact that analysts uncover information to which the econometrician is not privy. This information need not be "private", just unobserved by the econometrician (e.g., announcements by a drug company about doctors' willingness to prescribe a treatment), and hence not in the econometrician's model of valuation even when the information enters price. Moreover, the valuation consequences of such information surely persist—if an analyst has favorable information that the econometrician lacks, then some of that information likely remains months later. In our estimation, we must separate out and distinguish stickiness in recommendations due to an analyst's strategic design of recommendation "bins" and revision frictions from stickiness due

Figure 2: Recommendation cutoffs and revision frictions of a pretty good analyst

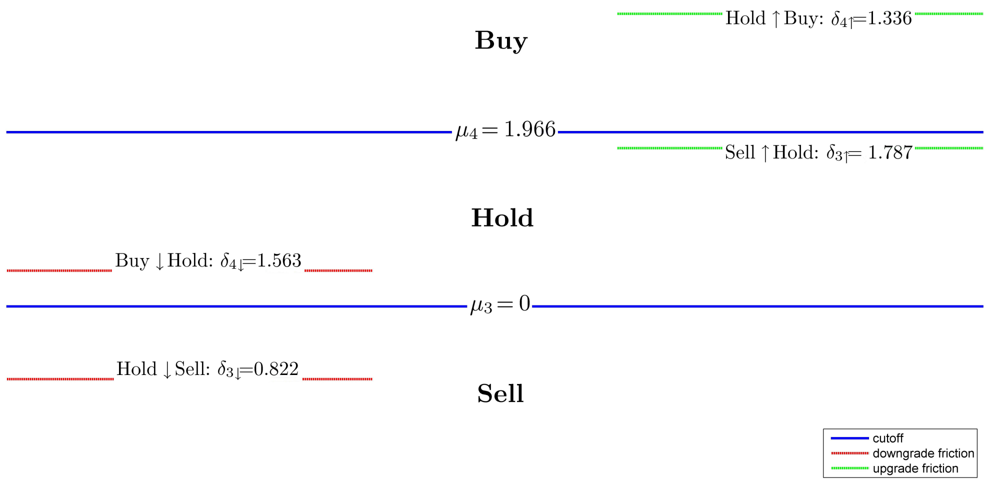

Cutoff-specific recommendation frictions, $\delta_{k\downarrow}$ and $\delta_{k\uparrow}$, bear the same index k as cutoff μ_k: $\delta_{k\downarrow}$ is the friction for downgrades *from* k to $k-1$, while $\delta_{k\uparrow}$ is the friction for upgrades from k *to* $k+1$.

to persistence in his "private information" (i.e., "information that the analyst has that the econometrician does not"). Since an analyst's assessment is not observable, we must treat these unobserved latent valuations as additional unknown parameters and analyze them jointly with other parameters using Monte Carlo methods.

To highlight the importance of posing the estimation in a dynamic framework, observe that an ordered probit approach in which recommendations are regressed on variables that capture components of value, estimates a model of an idealized analyst who lacks access to persistent "private" information (see Conrad et al. 2006). Our dynamic model allows us to address fundamental issues:

- How does an analyst choose the recommendation bin sizes, $\mu_{i+1} - \mu_i$, that determine the likelihoods of different initial recommendations, and how does he select the sizes of different recommendation revision frictions?

- An analyst can reduce the likelihood of frequent recommendation revisions not only with symmetric frictions, but also with asymmetric ones, where say the friction from buy to hold is large, but that from hold to buy is small. Since an analyst has flexibility in the design of frictions that limit recommendation revisions, how does he tailor them?

- Some brokerage houses employ five-tier rating systems, while others use a three-tier system (Buy, Hold or Sell). Do they design recommendations in similar ways?

- Our model has multiple sources that can raise the duration of a recommendation. The expected duration of a recommendation i increases in (a) its bin size, $\mu_{i+1} - \mu_i$, (b) its revision frictions, $\delta_{i+1\uparrow}$ and $\delta_{i\downarrow}$, and (c) persistence in "private" analyst information. What are the key drivers of stickiness for different recommendations?

- An analyst's notion of value may reflect not only factors that capture intrinsic asset fundamentals, but also attributes that appeal to his retail investor audience (e.g., small, growth stocks) and attributes such as underwriting business that matter to the analyst, but not his retail investor audience. Which firm characteristics positively affect an analyst's assessment, and which ones reduce it?

We test the model using analyst recommendations from the post Reg-FD,[1] post Global Analyst Research Settlement period, where analysts could issue negative sell recommendations without fear of losing access to company information sources. We estimate separate models for brokerage houses that employ three-tier rating systems and those that use the traditional five-tier rating system,[2] and for various subsamples (e.g., of larger brokerages).

The publicly-available characteristics that we find enter a pretty good analyst's assessment of value positively tend to be consistent with the findings of others (e.g., Conrad et al. (2006)). For example, a pretty good analyst has higher assessments of firms with (a) higher forecasted earnings or greater earnings surprises, (b) attributes that appeal to retail clients (small firms, as measured by market capitalization, book-to-price or sales growth; glamor firms, as measured by analyst coverage or institutional holdings), (c) less uncertainty (as measured by forecast dispersion), and (d) investment banking relationships with their brokerage house. An important exception is that, contrary to existing findings, once one controls for the pretty good analyst's strategic behavior and private information, measures of *past* firm performance (lagged returns) cease to positively affect assessments. That is, the idealized analyst model provides a misleading indication of the impacts of past firm performance.

[1]Reg FD was designed to curb the practice of selective disclosure of material nonpublic information. Reg FD eliminated incentives of analysts to issue favorable recommendations in order to curry favor with firms and hence retain access to nonpublic information. Reg FD came into effect in August 2000.

[2]Kadan et al. (2010) find that following the Global Analyst Research Settlement and related regulations on sell-side research in 2002, many brokerage houses, especially those affiliated with investment banks, switched from a five-tier rating system to a three-tier system; and subsequently more have switched.

What is fundamentally more important than these results is the fact that these readily-available public information sources matter *far less* for explaining the dynamics of analyst recommendations than do the recommendation frictions and the persistent analyst information. Highlighting this, the Bayes factor (the ratio of the marginal likelihoods of the alternative and null models) is an astonishing $exp(81660)$ for (a) the null model of an idealized analyst that includes all standard public information sources of valuation, but no persistent private information and (b) a barebones alternative model of a pretty good analyst that includes **no** public information components of valuation, and only two recommendation revision frictions, one for upgrades and one for downgrades, plus persistent analyst information.

The result that an analyst's "private information" matters more for recommendations than public information available to the econometrician is quite important. It suggests that most of the econometrician's information is already incorporated into prices, and hence has only secondary impacts on recommendations. Moreover, if recommendations only reflected readily-available information, they would have modest value, and one would be hard-pressed to justify why financial analysts should be well-paid. We find that about one-third of the valuation consequences of an analyst's "private" information persists to the next month.

The recommendation revision frictions that analysts introduce are as important for model fit as the analysts' information. Failing to account for these frictions biases up the estimate of the persistence in information, almost *tripling* the estimate. We find that analysts tailor revision frictions asymmetrically, depending on the recommendation. Analysts introduce much smaller frictions "out" of hold recommendations than "into" hold recommendations. This suggests that analysts do not like to maintain hold recommendations, presumably because they generate less trading volume for the brokerage house. For analysts using a five-tier rating system, we find that sell and buy recommendation bins are small relative to the frictions from strong sell to sell and strong buy to buy, so that most revisions are to hold. This, too, suggests strategic considerations: revisions from strong buy to buy that maintain a positive assessment or from strong sell to sell that maintain a negative assessment may not be enough to induce customers to unwind positions, but larger revisions to hold may do so.

We find that analysts who use the same (e.g., three-tier) recommendation rating system are well-described by a common model of recommendation bins and frictions, where sources of heterogeneity (firm and analyst attributes) only enter an analyst's valuation assessment. To verify the validity of this conclusion, we estimate our model on subsamples where one might suspect that analysts' recommendations might vary:

1. Over time. Estimates based on later years (2006-2010) are virtually identical to those for our entire sample. There is *no* evidence that analysts have altered how they issue recommendations over time. This means that differences between our estimates for three- and five-tier recommendation rating systems do not reflect temporal changes.

2. By brokerage size or analyst experience. Large brokerages ($>$ 52 analysts) and senior analysts ($\geq$ 5 years experience) have modestly higher estimates ($<$10% higher) of information persistence, and slightly higher frictions for revisions from hold to sell.

3. By analyst following. For heavily followed ($\geq$ 15 analysts) firms, information persistence is moderately ($<$ 25%) higher, but other structural parameters differ by little.

Overall, our estimates are remarkably robust: the homogeneous model of recommendation formation and revision describes analyst behavior extremely well.

As a final validation test, we investigate whether some of the stickiness that we find might proxy for analysts' imperfect and delayed reaction to new information (see Raedy, Shane and Yang (2006)): we estimate a model in which analysts may process a fraction of new information immediately, and the rest with a lag. Our estimates suggest slightly delayed incorporation of information, but they indicate that over 90% is processed immediately. Allowing for delayed incorporation of information *raises* the estimate of information persistence by roughly one-third, and has only modest (5-15%), non-uniformly signed impacts on recommendation frictions. This latter result reflects that absent non-trivial recommendation frictions, frequent revisions are likely whenever valuations are close to a recommendation cutoff.

Having validated the model structure, we investigate its indirect implications. Our model provides a theoretical lens through which to understand the different impacts of recommendation revisions made inside versus outside earnings announcement windows. Information arrival tends to be smooth between earnings announcements, but lumpy/discontinuous at announcements, or when firms provide earnings guidance. As a result, recommendation revisions made outside earnings announcement or guidance windows tend to occur when a value-price differential smoothly crosses a recommendation cutoff, whereas recommendation revisions made just after announcements often reflect discontinuous jumps in valuation information, so that value-price differentials "jump" past the revision cutoff. As a result, our model predicts that the duration of a revised recommendation "before retracement" to the original recommendation should be greater for recommendations made just after earnings announcements or guidance. Consistent with this, we find that the mean duration of a recommendation revision is 6-8 days longer if it is issued in a three-day window (on and after) of an earnings announcement or guidance date, than if it is issued outside these windows.

We also predict that recommendation revisions made inside earnings announcement or guidance windows should have greater impacts on stock prices than revisions made at other times. Consistent with this prediction, upgrades issued in the two days after earnings announcements or guidance result in three day cumulative abnormal returns that are 0.6 to 1.1 percentage points larger than those associated with revisions made at other times; and CAR differences for downgrades are *triple* those for upgrades. We then exploit the fact that the discontinuity in valuation assessment *only* occurs for revisions and *not* new recommendations. A difference-in-difference analysis of revisions vs. new recommendations inside and out of announcement and guidance windows provides strong reinforcing evidence for our model. Posed differently, our model reconciles the different market responses to new and revised recommendations made inside vs. out of earnings and guidance windows.

We conclude by exploiting the fact that our model provides a measure of the "surprise" associated with a recommendation revision or initiation. For example, if the current estimate of an analyst's stock valuation given *publicly-available* information is below an upgrade revision cutoff, the market should be more surprised by an upgrade, than if the estimated valuation suggests that the revision should already have been made. Concretely, the market should be more surprised by an upgrade to a buy if a stock's public information value suggests a hold than if it already suggests a buy. Thus, we predict a difference in the (appropriately signed) CARs following revisions in these two scenarios. So, too, when an analyst initiates coverage, the market response to a buy recommendation should be smaller if the current assessment of value given public information suggests a buy recommendation than if it indicates a hold.

We find both of these CAR relationships in the data. These findings provide strong confirmatory evidence for our model. In particular, they indicate that (a) investors believe financial analysts make recommendations along the lines of our model, and (b) the market values the information that analysts acquire that the econometrician does not possess. This means that the impacts of any unmodeled behavior-distorting incentives on recommendation formation, which serve to add noise, are modest enough that we still uncover these CAR relationships.

The paper is organized as follows. We next review the related literature. Section 2 develops our dynamic model of the analyst recommendation process. Section 3 provides an overview of our estimation approach. Section 4 details our data. Section 5 presents our findings. Section 6 concludes. An appendix includes more details on our estimation procedure.

1.1 The Literature

Our paper differs sharply from the existing literature in both its focus and methodology. Methodologically, most existing empirical analyses employ combinations of the following methods: (a) descriptive summary statistics, (b) investment strategies based on portfolio construction using recommendation information (Barber et al. (2001), Jegadeesh et al. (2004), Jegadeesh and Kim (2010), Boni and Womack (2006)), (c) regressions of recommendations on returns or regressions of returns on recommendations, and/or similar regressions based on *changes* of recommendations (Jegadeesh et al. (2004), Jegadeesh and Kim (2010), Barber et al. (2005), Ivkovic and Jegadeesh (2004), Bagnoli, Clement, Crawley, and Watts (2009)), and (d) correlation analysis (Boni and Womack (2006)).

The most closely-related paper is Conrad et al. (2006). They employ an ordered probit model coupled with a valuation model similar to ours to estimate a model of an idealized analyst who lacks access to persistent "private" information. This intrinsically static approach allows them to focus on subsamples surrounding large return events, and analyze the response of recommendations to major news. They find that analysts respond asymmetrically to positive and negative price shocks. Assuming away persistence in analyst information and recommendation frictions allows Conrad et al. (2006) to estimate their model using maximum likelihood. However, ordered probit models cannot deliver the intertemporal stickiness in recommendations found in the data, and we document the large biases that result. We extend their structural model of an idealized analyst by integrating persistent analyst information, thereby making the model explicitly dynamic. Despite a superior fit, this model still delivers far too little persistence in recommendations.[3]

Other studies provide descriptive statistical relationships between key analyst and firm characteristics and recommendations such as sample means, correlations, or quintiles (Jegadeesh et al. (2004), Ivkovic and Jegadeesh (2004)). Boni and Womack (2006) use a rank test for serial correlation between returns and recommendations. These raw empirical relationships provide valuable qualitative information about recommendations, offering useful guidance for building econometric models that can explain the complicated dynamics of the analyst recommendation process.

There is also a large literature focusing on the return reaction to recommendation changes,

[3]This stickiness in recommendations contrasts with the extreme "anti-herding" in forecasts of earnings that analysts issue to distinguish themselves. For example, the probability an analyst's forecast exceeds earnings given that it exceeds the consensus forecast is about 0.6 (Bernhardt, Campello and Kutsoati 2006).

and the investment value contained in recommendations. Stickel (1995) and Womack (1996) show that upgrades are associated with positive returns at the time of announcement. Womack (1996) finds that a post-recommendation price drift lasts up to one month for upgrades and six months for downgrades. Barber et al. (2001, 2005) find that absent transactions costs, investors could profit from the information in recommendations. Jegadeesh et al. (2004) investigate whether recommendations have investment value in the sense of predicting future returns, finding that analysts tend to issue more favorable recommendations for stocks with positive momentum and higher trading volume, and that analysts fail to respond quickly to negative signals by downgrading stocks. Jegadeesh and Kim (2006) confirm these findings in international markets. Ivkovic and Jegadeesh (2004) find a sharp increase in the information content of recommendation upgrades (but not downgrades) before earnings announcements.

Ljungqvist, Marston and Wilhelm (2006) use a limited dependent variable model to examine securities underwriting mandates and investigate the impact of investment banking relationship on recommendations. Boni and Womack (2006) study whether analysts chase or respond to price momentum at the industry level. They find that recommendation information is quite valuable for identifying short-term, within-industry mispricing. Loh and Stulz (2011) look at 2-day cumulative abnormal returns around recommendation changes, and study when recommendation changes are influential.

2 The Model

This section develops our dynamic model of the analyst recommendation process in the context of an analyst who employs a five-tier rating system; the model of an analyst who employs a three-tier rating system is similar.

If recommendations reflect buying opportunities, then they should reflect the difference between an analyst's assessment of a stock's valuation and its share price. This valuation assessment may reflect expected discounted earnings, or technical considerations that reflect market mispricing; and it may be prospective (e.g., the analyst's forecast of firm value in a year's time). Moreover, the notion of value is from the *analyst's* perspective: in addition to standard valuation fundamentals, attributes that appeal to an analyst's retail investor audience (e.g., small, growth, glamor stocks), or attributes such as underwriting business that only the analyst cares about, may enter the analyst's assessment of value.

Let V^*_{ijt} be analyst i's per share valuation of stock j at time t, which equals the per-share difference between the analyst's assessment of "value V_{ijt}" and price (P_{jt}):

$$V^*_{ijt} = \frac{V_{ijt} - P_{jt}}{P_{jt}}.$$

We assume that V^*_{ijt} is determined by a large set of explanatory variables that we describe later and analyst i's private information. Letting X_{ijt} be the per-share analogue of these variables, we write analyst i's per-share valuation model as:

$$V^*_{ijt} = X'_{ijt}\beta + u_{ijt}. \tag{1}$$

The unobserved residual terms u_{ijt} capture information that the analyst has, to which the econometrician is not privy. These residual terms are serially correlated: we must allow for persistence in u_{ijt} to capture the fact that the valuation consequences of this information will last for some time. Accordingly, we suppose that u_{ijt} evolves according to an AR(1) process:

$$u_{ijt} = \rho u_{ij,t-1} + \varepsilon_{ijt}, \tag{2}$$

where ε_{ijt} are i.i.d. $\mathrm{N}(0, \sigma^2)$ and ρ measures the persistence in the valuation consequences of the analyst's "private" information. For identification purposes, we normalize $\sigma^2 = 1$.[4]

As we highlighted in our introduction, an analyst i's recommendation for stock j at date t, R_{ijt}, is a function of both his valuation V^*_{ijt}, and his outstanding recommendation. Consequently, the model that determines an analyst's recommendation when he initiates coverage is *not* the same as the one that he uses to determine subsequent recommendations. In particular, an analyst's outstanding recommendation affects subsequent recommendations.

Suppose that analyst i initiates coverage for stock j at time t_{ij0}. His initial recommendation of $R_{ij,t_{ij0}}$ at $t = t_{ij0}$ is determined by the level of his valuation $V^*_{ijt_{ij0}}$ relative to the critical cutoffs $\mu_2 < \mu_3 < \mu_4 < \mu_5$ that he sets. Analyst i initiates coverage with a strong buy if $V^*_{ij,t_{ij0}} \geq \mu_5$, and with an appropriate lower recommendation if $V^*_{ij,t_{ij0}}$ falls into the corresponding valuation bin. That is,

$$R_{ij,t_{ij0}} = \begin{cases} 5, & \text{if } \mu_5 \leq V^*_{ijt_{ij0}} \\ 4, & \text{if } \mu_4 \leq V^*_{ijt_{ij0}} < \mu_5 \\ 3, & \text{if } \mu_3 \leq V^*_{ijt_{ij0}} < \mu_4 \\ 2, & \text{if } \mu_2 \leq V^*_{ijt_{ij0}} < \mu_3 \\ 1, & \text{if } \qquad V^*_{ijt_{ij0}} < \mu_2. \end{cases} \tag{3}$$

Thus, a recommendation of a 5 represents a strong buy, a 4 is a buy, a 3 is a hold, a 2 is a sell, and a 1 is a strong sell. Without loss of generality, for identification purposes, we normalize μ_2 to zero.

[4]The variance just multiplicatively scales the cutoffs for recommendations and revisions.

Subsequent recommendations R_{ijt} are determined by both the analyst's updated valuation V^*_{ijt} and his outstanding recommendation $R_{ij,t-1}$. Our model captures an analyst's reluctance to change recommendations via recommendation-specific revision frictions that the analyst can introduce. Specifically, if analyst i's outstanding recommendation for stock j at time t is $R_{ij,t-1} = k$, then analyst i would not lower his recommendation R_{ijt} to $k-1$ unless his valuation V^*_{ijt} falls below the threshold value μ_k by an amount $\delta_{k\downarrow}$ that the analyst chooses. Similarly, analyst i will not raise his recommendation R_{ijt} to $k+1$ unless $V^*_{ijt} > \mu_{k+1} + \delta_{k+1,\uparrow}$. In other words, the analyst effectively expands the bin corresponding to his previous recommendation k from $[\mu_k, \mu_{k+1})$ to $[\mu_k - \delta_{k\downarrow}, \mu_{k+1} + \delta_{k+1,\uparrow})$, and does not revise his recommendation unless his valuation assessment V^*_{ijt} evolves outside of this expanded bin. See Figure 2. Such revision frictions are "localized" in that (a) the extents to which a recommendation bin is expanded can depend on the recommendation itself (i.e., $\delta_{k+1,\uparrow}$ and $\delta_{k\downarrow}$ can vary with k), and (b) revision frictions only affect decisions to upgrade or downgrade to "adjacent" recommendations. For example, if an analyst has an outstanding strong sell ($R = 1$) rating, the friction $\delta_{2\uparrow}$ only affects upgrades to a sell: as long as $\mu_3 > \mu_2 + \delta_{2\uparrow}$, this friction has no effect on upgrades to hold. Therefore, the probability distribution over analyst i's recommendations for stock j at time t is

$$\Pr\left[R_{ijt} = k | X_{ijt}, R_{ij,t-1}\right] = \begin{cases} \Pr\left[\mu_k \leq V^*_{ijt} < \mu_{k+1}\right], & \text{if } R_{ij,t-1} > k+1, \\ \Pr\left[\mu_k \leq V^*_{ijt} < \mu_{k+1} - \delta_{k+1,\downarrow}\right], & \text{if } R_{ij,t-1} = k+1, \\ \Pr\left[\mu_k - \delta_{k\downarrow} \leq V^*_{ijt} < \mu_{k+1} + \delta_{k+1,\uparrow}\right], & \text{if } R_{ij,t-1} = k, \\ \Pr\left[\mu_k + \delta_{k\uparrow} \leq V^*_{ijt} < \mu_{k+1}\right], & \text{if } R_{ij,t-1} = k-1, \\ \Pr\left[\mu_k \leq V^*_{ijt} < \mu_{k+1}\right], & \text{if } R_{ij,t-1} < k-1, \end{cases} \tag{4}$$

where we adopt the convention that $\mu_1 = -\infty$ and $\mu_6 = +\infty$. This formulation allows for distinct recommendation-specific frictions for both upgrades and downgrades, for example, for $\delta_{k\uparrow} \neq \delta_{k\downarrow}$. This model nests several more restricted models:

- If $\delta_{k\uparrow} = \delta_\uparrow$ and $\delta_{k\downarrow} = \delta_\downarrow$ for each k, then the upgrade friction can differ from the downgrade friction, but these two frictions do not depend on the recommendation level.

- $\delta_{k\uparrow} = \delta_{k'\downarrow} = 0$, for all k, k', captures the "idealized" financial analyst who implicitly has been the focus of existing literature.

Equations (1) – (4) lay out the model's econometric structure. Equations (1) and (2) capture the dynamics of analyst stock valuations, while equations (3) and (4) govern analyst decisions on initial recommendations, and subsequent revisions and reiterations.

3 Overview of Model Estimation

This section presents the parameter estimation procedure, and the metrics used to evaluate the goodness of fit of alternative model specifications. The appendix provides more details.

Parameter Estimation. In practice, we observe X_{ijt}, but not V_{ijt}^*. That is, V_{ijt}^* is a latent variable. In addition, u_{ijt} are serially correlated over t. Our model is even more complicated than an ordered probit model with serial correlation, as the recommendation revision friction parameters make the estimation even more nonlinear, adding challenge to an already complicated estimation problem. This structure creates barriers to estimating the model using traditional MLE or GMM methods. For the reasons discussed in the literature on estimating dynamic ordered probit models, simulation-based methods are typically used to estimate this class of models. Accordingly, we estimate our model using Markov Chain Monte Carlo (MCMC) methods. The MCMC approach is conceptually easy to implement here because we can partition the parameter space in such a way that the conditional posterior densities take a simple form, making it tractable to draw random variables from their conditional distributions. Once we reach the stationary distribution for the parameter estimates, we routinely obtain parameter estimates from their sample averages over 150,000 draws.

Consider first the simple case where analyst i begins giving recommendations for firm j at time t_{ij0}, and continues until time T. We denote the random sequence of recommendation choices by analyst i for firm j by:

$$\mathbf{R}_{ij} = (R_{ijt_{ij0}}, \cdots\cdots, R_{ijT})'$$

and observations (realizations) of $\mathbf{R}_{ij}$ by

$$\mathbf{r}_{ij} = (r_{ijt_{ij0}}, \cdots\cdots, r_{ijT})',$$

where $r_{ijt} \in [1, 2, 3, 4, 5]$. We denote the vector of unknown parameters by $\theta = (\beta', \mu', \delta', \rho)'$. Then, denoting the information set up to time t by $\mathcal{F}_{t-1}$ (i.e., $\mathcal{F}_{t-1}$ contains information about previous recommendations, updated public information, etc.), given the parameters θ and conditional on the information at the starting date, we have[5]

$$P\left[\mathbf{R}_{ij} = \mathbf{r}_{ij}\right] = P\left[R_{ijt_{ij0}} = r_{ijt_{ij0}}\right] \prod_{t=t_{ij0}+1}^{T} P\left[R_{ijt} = r_{ijt} | \mathcal{F}_{t-1}\right].$$

[5]Throughout, we consider distributions conditional on available information at the date coverage is initiated. To ease presentation, we omit this conditioning in our notation. For example, Pr $\left[R_{ijt_{ij0}} = r_{ijt_{ij0}}\right]$ denotes the distribution over the initial recommendation conditional on information available then.

Given the information at the starting date when coverage is initiated, the probability distribution over initial recommendations is determined by (3), so that

$$P\left[R_{ijt_{ij0}} = r_{ijt_{ij0}}\right] = P\left[\mu_{r_{ijt_{ij0}}} \leq V^*_{ijt_{ij0}} < \mu_{r_{ijt_{ij0}}+1}\right].$$

Because the list of variables used when coverage is initiated differs from later dates (lagged recommendations do not enter when coverage is initiated), we denote the vector of right-hand side variables when coverage is initiated by $\underline{X}_{ij,t_{ij0}}$. The analyst's initial valuation model is

$$V^*_{ij,t_{ij0}} = \underline{X}'_{ij,t_{ij0}}\underline{\beta} + u_{ij,t_{ij0}}.$$

Therefore, from the AR structure of the analyst's information process, equation (2), we have

$$P[R_{ijt_{ij0}} = r_{ijt_{ij0}}] = \Phi\big(\sqrt{1-\rho^2}[\mu_{k_{t_{ij0}}+1} - \underline{X}'_{ij,t_{ij0}}\underline{\beta}]\big) - \Phi\big(\sqrt{1-\rho^2}[\mu_{k_{t_{ij0}}} - \underline{X}'_{ijt_{ij0}}\underline{\beta}]\big). \quad (5)$$

Once coverage has been initiated, the conditional distributions, $P\left[R_{ijt} = r_{ijt}|\mathcal{F}_{t-1}\right]$, are determined by (2) and (4) together. Let the vector of right-hand side variables after coverage has been initiated be X_{ijt}. Then $V^*_{ijt} = X'_{ijt}\beta + u_{ijt}$, and defining

$$g_{ijt}(\theta) = \rho V^*_{ij,t-1} + \left(X'_{ijt} - \rho X'_{ij,t-1}\right)\beta,$$

we have

$$\begin{aligned} P\left[R_{ijt} = r_{ijt}|\mathcal{F}_{t-1}\right] \;=\; & 1(r_{ijt} = r_{ij,t-1})\left[\,\Phi\left(\mu_{r_{ijt}+1} - g_{ijt}(\theta) + \delta_{r_{ijt}+1,\uparrow}\right) - \Phi\left(\mu_{r_{ijt}} - g_{ijt}(\theta) - \delta_{r_{ijt},\downarrow}\right)\right] \\ & + 1(r_{ijt} = r_{ij,t-1} - 1)\left[\,\Phi\left(\mu_{r_{ijt}+1} - g_{ijt}(\theta) - \delta_{r_{ijt}+1,\downarrow}\right) - \Phi\left(\mu_{r_{ijt}} - g_{ijt}(\theta)\right)\right] \\ & + 1(r_{ijt} < r_{ij,t-1} - 1)\left[\,\Phi\left(\mu_{r_{ijt}+1} - g_{ijt}(\theta)\right) - \Phi\left(\mu_{r_{ijt}} - g_{ijt}(\theta)\right)\right] \\ & + 1(r_{ijt} = r_{ij,t-1} + 1)\left[\,\Phi\left(\mu_{r_{ijt}+1} - g_{ijt}(\theta)\right) - \Phi\left(\mu_{r_{ijt}} - g_{ijt}(\theta) + \delta_{r_{ijt},\uparrow}\right)\right] \\ & + 1(r_{ijt} > r_{ij,t-1} + 1)\left[\,\Phi\left(\mu_{r_{ijt}+1} - g_{ijt}(\theta)\right) - \Phi\left(\mu_{r_{ijt}} - g_{ijt}(\theta)\right)\right]. \end{aligned} \quad (6)$$

Thus, letting $\pi(\theta)$ be the prior, we can write the joint distribution of data and parameters as

$$\pi(\theta, r) = \pi(\theta)\prod_{j=1}^{J}\prod_{i\in I_j}\left\{P(r_{ijt_{ij0}})\prod_{t=t_{ij0}+1}^{T} P(r_{ijt}|r_{ij,t-1})\right\},$$

where r is the vector of all realizations of recommendations, $P(r_{ijt_{ij0}}) = P(R_{ijt_{ij0}} = r_{ijt_{ij0}})$ and $P(r_{ijt}|r_{ij,t-1}) = P\left[R_{ijt} = r_{ijt}|\mathcal{F}_{t-1}\right]$ are defined by (5) and (6).

Analyst recommendations and related firm- and analyst-level control variables represent an unbalanced panel dataset containing observations of many analyst-firm pairs (a particular firm followed by a particular analyst) over multiple time periods, reflecting that an analyst may cease following a stock for some time, but then re-initiate coverage. In this case, analyst i may issue recommendations for firm j in n_{ij} different periods, say $\{t_{ij0}(s), \cdots, t_{ij*}(s),\ s = 1, \cdots, n_{ij}\}$, during $t = 1, ..., T$. If we let

$$H_{ijs}\left(r_{ij}, \theta\right) = \left\{ P(r_{ijt_{ij0}(s)}) \prod_{t=t_{ij0}(s)+1}^{t_{ij*}(s)} P(r_{ijt}|r_{ij,t-1}) \right\},$$

where the probability $P(r_{ijt_{ij0}(s)})$ is defined by (5) and conditional probability $P(r_{ijt}|r_{ij,t-1})$ is given by (6), then the joint distribution of data and parameters is

$$\pi(\theta, r) = \pi(\theta) \prod_{j=1}^{J} \prod_{i \in I_j} \prod_{s=1}^{n_{ij}} H_{ijs}\left(r_{ij}, \theta\right).$$

In our Bayesian estimation approach, we treat the unobserved (latent) valuations as additional unknown parameters and analyze them jointly with the other parameters (θ) using Markov Chain Monte Carlo (MCMC) methods. Let R denote the observed analyst recommendations, and V denote the latent analyst valuations. We divide the vector of parameters θ into 4 groups: (1) valuation parameters β; (2) recommendation bin parameters μ_j, $j = 3, 4, 5$; (3) recommendation revision friction parameters δ (i.e., $\delta_{k\uparrow}, \delta_{k\downarrow}$, etc); and (4) the information persistence parameter ρ.

The MCMC estimator using a Gibbs sampler starts with an initial value $(\theta^{(0)}, V^{(0)})$, and then simulates in turn. The above partition of the parameter vector yields a reasonably simple form for the conditional posterior densities that makes it practically easy to draw random variables from the conditional distributions. The distribution of β conditional on the data and other parameters is normal with mean

$$\begin{aligned} \widehat{\beta} &= \left[\sum_{j=1}^{J} \sum_{i \in I_j} \prod_{s=1}^{n_{ij}} \prod_{t=t_{ij0}(s)+1}^{t_{ij*}(s)} \left(X_{ijt} - \rho X_{ij,t-1}\right) \left(X'_{ijt} - \rho X'_{ij,t-1}\right) \right]^{-1} \\ & \sum_{j=1}^{J} \sum_{i \in I_j} \prod_{s=1}^{n_{ij}} \prod_{t=t_{ij0}(s)+1}^{t_{ij*}(s)} \left(X_{ijt} - \rho X_{ij,t-1}\right) \left(V^*_{ijt} - \rho V^*_{ij,t-1}\right), \end{aligned}$$

and variance

$$\Sigma_{\widehat{\beta}} = \left[\sum_{j=1}^{J} \sum_{i \in I_j} \prod_{s=1}^{n_{ij}} \prod_{t=t_{ij0}(s)+1}^{t_{ij*}(s)} \left(X_{ijt} - \rho X_{ij,t-1}\right) \left(X'_{ijt} - \rho X'_{ij,t-1}\right) \right]^{-1},$$

where $X_{ij0} = 0$. The conditional distribution of the persistence parameter ρ is a truncated normal distribution, where the truncation reflects $|\rho| \in (0, 1)$. The conditional distributions of the δ recommendation friction and μ bin parameters are uniform distributions, where the bounds that reflect that given V_{ijt} and other parameters, the parameter must be consistent with the realized recommendation r_{ijt}. Lastly, the conditional distributions of the latent variables V^*_{ijt} are truncated normal distributions, where the truncation reflects consistency of recommendation r_{ijt} with the frictions and bin parameters.

After each draw of a new value of a parameter, the corresponding subvector of previous values is replaced by the new subvector that has the new value rather than the old one. We then continue to draw a new value of another parameter. Thus, on the q^{th} draw,

1. We first draw $V^{(q)}$ from the conditional distribution $p(V|\beta^{(q-1)}; \mu_j^{(q-1)}; \delta^{(q-1)}; \rho^{(q-1)})$
2. Update $V^{(q)}$ and draw $\beta^{(q)}$ from the conditional distribution $p(\beta|V^{(q)}; \mu_j^{(q-1)}; \delta^{(q-1)}; \rho^{(q-1)})$
3. Draw $\mu_j^{(q)}$ $j = 3, 4, 5$, from the conditional distribution $p(\mu|V^{(q)}; \beta^{(q)}; \delta^{(q-1)}; \rho^{(q-1)})$:
   (a) draw $\mu_3^{(q)}$ from the conditional distribution $p(\mu_3|V^{(q)}; \beta^{(q)}; \mu_4^{(q-1)}, \mu_5^{(q-1)}; \delta^{(q-1)}; \rho^{(q-1)})$.
   (b) draw $\mu_4^{(q)}$ from the conditional distribution $p(\mu_4|V^{(q)}; \beta^{(q)}; \mu_3^{(q)}, \mu_5^{(q-1)}; \delta^{(q-1)}; \rho^{(q-1)})$.
   (c) draw $\mu_5^{(q)}$ from the conditional distribution $p(\mu_5|V^{(q)}; \beta^{(q)}; \mu_3^{(q)}, \mu_4^{(q)}; \delta^{(q-1)}; \rho^{(q-1)})$.
4. Sequentially draw $\delta^{(q)}$'s and $\rho^{(q)}$ in the same way.
5. Set $q = q + 1$, and repeat.

After a complete iteration, we obtain an updated vector $(\theta^{(q)}, V^{(q)})$. We repeat this P times, and as $P \to \infty$, the distribution of $(\theta^{(P)}, V^{(P)})$ converges to the distribution of (θ, V).

Once we have a set of reasonable initial values, to ensure that we sample from the stationary distribution for the parameter estimates, we discard the initial $50,000$ iterations,[6] and keep the next $150,000$ iterations as sample draws. To test the null hypothesis that the Markov chain of each parameter estimate is from the stationary distribution, we use Geweke's (1992) convergence diagnostic, which tests that the means of the first 10% of the Markov chain (observations $50,001 - 65,000$) and the last 50% of observations are equal. The data easily pass

[6]Prior to identifying a set of reasonable starting values, we discarded as many as 200,000 observations.

this convergence check. We also visually check the trace plots of each Markov chain to confirm that it has a relatively constant mean and variance. Parameter estimates are then given by their sample averages. The appendix provides more details on estimation procedures.

The computation time for each model is about 7 days. The computational burden is mainly due to the slow convergence of the Monte Carlo Markov chains for the cutoff estimates. Intuitively, the slow convergence in the cutoff estimates μ_k is due to the narrow interval of its conditional uniform distribution, which leaves little room for a cutoff estimate to move toward its true value within one iteration. Similar findings have been reported for estimation of the independent multinomial model (Cowles (1996)) and autoregressive ordered probit model (Muller and Czado (2005)).

Goodness of Fit. We use the Brier Score (Brier, 1950) and Bayes factor to compare the goodness of fit of different model specifications. The Brier score is the mean squared deviation between the observed outcome and the (in-sample) predicted probability of a recommendation:

$$S = \frac{1}{\#obs} \sum_{ijt} \left[I(R_{ijt}) - \widehat{\pi}(R_{ijt}) \right]^2 ,$$

where $I(R_{ijt})$ is an indicator function, equal to one if recommendation R_{ijt} is observed, and zero otherwise; $\widehat{\pi}(R_{ijt})$ denotes the *in-sample estimate* of the probability of recommendation R_{ijt}; and $\#obs$ is the total number of recommendations issued in our sample. The Brier score penalizes large deviations in a probability forecast: the smaller is S, the better is the model fit. A perfect probability forecast would yield a Brier score of zero.

More formally, we employ the Bayes factor to assess the goodness of model fit. All specifications considered in our study are nested in the full model. We have no prior belief over the null model and the alternative (i.e., $\Pr(M_0) = \Pr(M_A) = 0.5$). Given the observed data, D, the Bayes factor, B, is defined as

$$B = \frac{\Pr(D|M_A)}{\Pr(D|M_0)},$$

where $\Pr(D|M_0)$ and $\Pr(D|M_A)$ are the marginal likelihoods of the null and alternative models, respectively. In terms of the logarithm of models' marginal likelihoods (reported in Table 5), the Bayes factor is $\exp(\log(\Pr(D|M_A)) - \log(\Pr(D|M_0)))$. Kass and Raftery (1995) argue that a Bayes factor of $2\log(B)$ that exceeds 10 ($\approx B > 150$) represents decisive evidence in favor of the alternative model against the null. A Bayes factor that exceeds $1,000$ ($B > 1000$) provides conclusive support for forensic evidence in a criminal trial (Evett 1991). The Bayes factor also penalizes overfitting (over-parametrization) of an alternative model.

4 Data

Our sample of analyst recommendations is from the Institutional Brokers Estimate System (I/B/E/S) U.S. Detail file. Following most studies in the literature, we reverse the I/B/E/S recommendation coding so that more favorable recommendations correspond to larger numbers (i.e., 1=Strong Sell, 2=Sell, 3=Hold, 4=Buy and 5=Strong Buy). Each analyst is identified by I/B/E/S/ with a unique numerical code (analyst masked code). We use this numerical identifier to match an analyst's stock recommendations to his earnings forecasts in the I/B/E/S Detail History file. We exclude recommendations issued by unidentified/anonymous analysts. Stock return and trading volume related data are collected from CRSP. Firms' accounting and balance-sheet information is extracted quarterly from Compustat.

We use monthly data from January 2003 to December 2010 from the post Regulation Fair Disclosure, post Global Analyst Research Settlement period. If an analyst issues multiple recommendations for a firm within a calendar month, we only use the last recommendation. Our choice of monthly frequencies reflects several practical considerations. First, analysts rarely change recommendations multiple times in a month (only 1.2% of recommendations in our sample are revised multiple times in a month) and there is a concern that analysts introduce slight temporal revision frictions, to avoid repeated revisions over a short period of time. It is not feasible to allow for temporal frictions in estimation; and monthly observations minimize the impact of any temporal frictions on estimates. One might also worry that information arrival is uneven due to endogenous information acquisition—an analyst who gathers information about a firm today is less likely to do so tomorrow, leading to lumpiness in information arrival at high frequencies, in essence, a high frequency source of stickiness. However, an analyst will monitor a firm more than once a month, so endogenous information acquisition should not lead to lumpy information arrival at monthly frequencies. Second, much of our data is observed at lower frequencies (e.g., sales growth or earnings). Third, monthly frequencies facilitate estimation, as the median time to recommendation revision is 190 days.

Brokerage houses that use three-tier rating systems appear in the I/B/E/S database as issuing either Buy, Hold and Sell (4, 3, 2) recommendations only; or as issuing only Strong Buy, Hold and Strong Sell (5, 3, 1) recommendations. We pool these two populations into a single three-tier rating system.[7] We identify the date at which brokerage houses switch to the three-tier system by the date at which they exclusively issue from that subset (they typically

[7] Estimates if we do not pool are qualitatively identical.

switch on the same day). Our sample of five-tier brokerage houses only includes those that never switch; and we only use observations on three-tier brokerage houses once they switch.

When analysts maintain a recommendation from one month to the next, they typically do not reiterate their recommendations. Moreover, analysts sometimes cease following a stock without indicating a stopped coverage on I/B/E/S. To avoid building in spurious persistence in estimates of private information and larger recommendation revision frictions by including non-varying recommendations from analysts who ceased following a stock, we conservatively assume that an analyst who does not reiterate or revise a recommendation within 12 months has dropped coverage. Thus, in the absence of a revision or reiteration or stopped coverage indication by analyst i for stock j (an analyst-firm pair) in month t, we set the recommendation, R_{ijt}, to be the most recent recommendation/reiteration issued in the past 12 months by the analyst for that firm.[8] For a given analyst-firm pair, an observed recommendation with no preceding outstanding recommendation in the past year is classified as an initiation, and a recommendation revision/reiteration refers to a recommendation for which there was an outstanding recommendation[9] in the previous month. We exclude analyst-firm pairs with fewer than 20 recommendations (including filled in reiterations) over the entire sample period. This policy largely eliminates only analysts who *never* revise or reiterate a recommendation, dropping analysts who may have quickly lost interest and ceased following a firm. Our final sample of analysts using a three-tier rating system consists of an unbalanced panel data with 241076 recommendations by 1927 analysts (from 188 brokerage houses) for 2805 firms (8224 analyst-firm pairs); and for analysts using a five-tier rating system, we have 89726 recommendations by 740 analysts (from 128 brokerage houses) for 1894 firms (3059 analyst-firm pairs).

Table 1 presents the distributions of recommendation levels and the transition matrix of recommendation revisions and reiterations for brokerage houses using a three-tier rating system; and Table 2 does so for those using a five-tier rating system. Almost half of the recommendations by brokerages that use three ratings are holds, 41% are buys and 10% are sells. Table 2 shows that brokerage houses that use five ratings are more optimistic—about 53% of their recommendations are strong buys or buys, and only 7% are sells or strong sells—likely reflecting that five-tier brokerage houses, which tend to be smaller, without an investment bank side, have different audiences. This means that one cannot collapse five-tier brokerage

[8] We show that our qualitative empirical findings are robust if we use different cutoffs (e.g., 9 or 15 months) to identify analysts who dropped coverage.

[9] This outstanding recommendation may be an actual issuance by the analyst or a carryover from a recent issuance within the past twelve months.

Table 1: Distribution of Analyst Recommendations (three-tier ratings)

Panel A. Stock Recommendation Levels				
	Buy, 4	Hold,3	Sell, 2	Total
Initiations	29910	35862	7172	72944
	41.00%	49.16%	9.83%	100%
Full Sample	99159	118626	23291	241076
	41.13%	49.21%	9.66%	100%

Panel B. Transition Matrix of Recommendation Revisions and Reiterations

	To: Buy, 4	Hold, 3	Sell, 2	Total
From:				
Buy, 4	65859	3314	100	69273
	95.07%	4.78%	0.14%	100%
Hold, 3	3265	78319	1026	82610
	3.95%	94.81%	1.24%	100%
Sell, 2	121	1127	14993	16241
	0.74%	6.94%	92.31%	100%
Total	69245	82760	16119	

houses into three-tier ones by grouping strong buys with buys, and strong sells with sells.

For brokerage houses using a three-tier system, transitions out of buy are about as likely as those out of hold, while upward transitions out of sell are about 50% more likely. Brokerage houses using a five-tier system do not hold negative ratings for as long as those using a three-tier system—they are more likely to revise holds or sells upward, and less likely to revise buy/strong buy ratings down—additional indications that they tailor recommendations more optimistically. Of note, brokerage houses using a five-tier system are more likely to revise recommendations to hold than to other revisions, *even* from strong buy and strong sell.

Public Information Components of Value. We consider a wide range of public information firm and analyst characteristics that plausibly enter an analyst's assessment of value, most of which have been suggested by prior studies to be related to recommendations.

We construct several variables from CRSP. Omitting analyst and firm subscripts (i and j), and the subscript t for the *calendar* month in which recommendation R_{ijt} is issued, we use:

ret_{-1}: stock excess return (on the market) in the past month (month $t-1$).

Table 2: Distribution of Analyst Recommendations (five-tier ratings)

Panel A. Stock Recommendation Levels

	Strong Buy, 5	Buy, 4	Hold,3	Sell, 2	Strong Sell, 1	Total
Initiations	5476	8137	9960	1519	364	25456
	21.51%	31.96%	39.13%	5.97%	1.43%	100%
Full Sample	19371	28816	35228	5068	1243	89726
	21.59%	32.12%	39.26%	5.65%	1.39%	100%

Panel B. Transition Matrix of Recommendation Revisions and Reiterations

To:	Strong Buy, 5	Buy, 4	Hold,3	Sell, 2	Strong Sell, 1	Total
From:						
Strong Buy, 5	12911	395	499	9	2	13816
	93.45%	2.86%	3.61%	0.07%	0.01%	100%
Buy, 4	498	19355	900	58	12	20823
	2.39%	92.95%	4.32%	0.28%	0.06%	100%
Hold,3	474	869	23512	251	60	25166
	1.88%	3.45%	93.43%	1.00%	0.24%	100%
Sell, 2	6	51	299	3220	23	3599
	0.17%	1.42%	8.31%	89.47%	0.64%	100%
Strong Sell, 1	6	9	58	11	782	866
	0.69%	1.04%	6.70%	1.27%	90.30%	100%
Total	13895	20679	25268	3549	879	

$ret_{-2:-6}$: median-term monthly excess return.

$ret_{-7:-12}$: long-term monthly excess return.

To preclude the impact of analysts' recommendations on stock performance (reverse-causality), returns are calculated as the holding period return based on monthly closing stock prices. For instance, $ret_{-2:-6}$ is the return received from holding a stock at the closing price in month $t-7$ and selling it at the closing price in month $t-2$.

$\sigma_{-1:-6}$: past six-month stock return volatility (daily return volatility times the square root of the number of trading days over the past six months).

$\log(turnover_{-1:-6})$: log of (trading volume in past 6 months scaled by shares outstanding).

Table 3: Descriptive Statistics for Firm and Analyst-related Variables

	Mean	Std. Dev.	$Quar_1$	$Quar_2$	$Quar_3$	Buy−Sell
ret_{-1}	0.750	11.078	−5.405	0.746	6.879	0.486***
$ret_{-2:-6}$	3.813	22.913	−9.584	3.268	15.882	2.602***
$ret_{-7:-12}$	4.492	26.689	−10.925	3.745	18.316	6.265***
$\sigma_{-1:-6}$	25.554	13.360	15.862	22.242	31.352	−2.212***
$\log(turnover_{-1:-6})$	−0.164	0.735	−0.641	−0.137	0.344	−0.041***
$\log(MktCap_{-1:-6})$	14.413	1.552	13.271	14.306	15.459	−0.104***
$\log(num\_anal)$	2.644	0.689	2.197	2.708	3.135	0.058***
$HSize$	3.742	1.113	3.178	3.951	4.673	−0.226***
SUE	0.330	1.451	−0.692	0.359	1.553	0.426***
BM	0.516	0.345	0.272	0.443	0.673	−0.098***
EP	0.017	0.109	0.018	0.043	0.064	0.025***
SG	1.115	0.212	0.999	1.089	1.199	0.057***
ROA	0.062	0.112	0.015	0.062	0.123	0.026***
$FAge$	22.364	18.774	9	16	32	−0.286**
$FRtoP$ ($\times 10^{-3}$)	−0.652	12.586	−2.951	0.605	3.435	2.489***
$CFtoP$ ($\times 10^{-2}$)	1.005	1.548	0.649	1.209	1.738	0.362***
$FDisp$ ($\times 10^{-2}$)	0.195	0.307	0.038	0.083	0.201	−0.096***
$FDev$ ($\times 10^{-3}$)	−0.103	4.898	−1.056	0	1.107	0.612***
IH	0.747	0.241	0.609	0.790	0.919	0.032***
$\log(year\_brkg)$	1.487	0.677	1.099	1.609	1.946	−0.091***
$\log(year\_IBES)$	1.893	0.589	1.386	1.946	2.398	−0.029***

Results of univariate comparisons of means between stocks with a Sell (Sell or Strong Sell) recommendations and those of Buy (Buy and Strong Buy) are reported in the far right column. *** and ** denote statistical significance at 1% and 5% levels.

$\log(MktCap_{-1:-6})$: log of a firm's market capitalization (monthly closing price times shares outstanding).

$Firm\_Age$: years a firm has been in the CRSP database at the calendar year of month t.

Firm accounting variables come from Compustat. Here q denotes the most recent *fiscal* quarter for which an earnings announcement was made prior to or within a calendar month t.

SUE: standardized unexpected earnings. $SUE = (EPS_q - EPS_{q-4})/std(EPS_{q:q-7})$, where $EPS_q - EPS_{q-4}$ is a firm's unexpected quarterly earnings per share and $std(EPS_{q:q-7})$ is the firm's earnings volatility over the eight preceding quarters. We require a firm to report EPS at least four times in the past eight quarters to calculate earnings volatility. The value of SUE is carried over the following quarter after the release of EPS_q.

$SUE \times D_{EA}$: interaction between an earnings-announcement-month dummy and standardized unexpected earnings. The D_{EA} dummy indicates an earnings announcement was made either in that month, or in the last five trading days of the previous month (to give the market time to assess the earnings).

BM: ratio of book equity to market equity in quarter q;

EP: earnings-to-price ratio. $EP = \sum_{i=0}^{3}(EPS_{q-i})/Prc_q$, where Prc_q is the stock price at the end of quarter q.

SG: annual sales growth rate. $SG = \sum_{i=0}^{3} Sales_{q-i} / \sum_{i=0}^{3} Sales_{q-4-i}$, where $Sales$ is a firm's quarterly total sales.

ROA: return on assets. $ROA = \sum_{i=0}^{3} Income_{q-i}/AT_q$, where $Income$ and AT are quarterly net income and the end-of-quarter total assets.

The I/B/E/S recommendation file yields the following analyst-related variables:

$\log(num\_anal)$: logarithm of the number of analysts with outstanding stock recommendations on a firm in month $t-1$.

$HSize$: logarithm of the number of analysts at a brokerage issuing stock recommendations over the course of one calendar year, as in Agrawal and Chen (2008).

We use the I/B/E/S Detail History and Summary Statistics files to construct:

$FRtoP$: Earnings forecast revisions to price ratio is the rolling sum of the preceding six months revisions to price ratios (Jegadeesh et al., 2004). $FRtoP = \sum_{i=1}^{6}(f_{t-i} - f_{t-1-i})/Prc_{t-1-i}$, where f_t is the mean consensus analyst quarterly forecast in month t.

$CFtoP$: Consensus quarterly earnings forecast to price ratio. $CFtoP = f_{t-1}/Prc_{t-1}$.

$FDisp$: Forecast dispersion is the standard deviation of analysts' quarterly earnings forecasts at month $t-1$ scaled by Prc_{t-1}.

$FDev$: Analyst earnings forecast deviation is the difference between an analyst's forecast and the consensus earnings forecast at month $t-1$ scaled by Prc_{t-1}.

Finally, we consider:

IH: Institutional Holdings is the percentage of a firm's equity held by institutional investors in quarter q, obtained from 13f quarterly filings to the Securities and Exchange Commission (Thomson Financial 13f institutional database).

D_{IB}: Dummy variable indicating an investment banking relationship (lead- or joint-management appointment) in the previous five years between the analyst's brokerage house and the firm. We obtain all US debt and equity offerings from the Security Data Company (SDC) Database.

$\log(year\_brkg)$: Logarithm of the years an analyst has been at his/her current brokerage firm.

$\log(year\_IBES)$: Logarithm of the years an analyst has been in the I/B/E/S database.

In addition to these 22 explanatory variables, we control for industry fixed effects captured by one-digit SIC codes. Table 3 provides summary statistics for these public information variables in the three-tier sample. The right-most column reports results of univariate comparisons of means between stocks with Sell and Buy recommendations. All differences in control variables between sell and buy recommendations are significant at a 5% significance level.

5 Empirical Analysis

We next present estimates of our dynamic model of analyst recommendations. We compare results from the full model (detailed in equations (1)–(4)) with those from more restricted models to emphasize the importance of both information persistence and revision frictions in the analyst decision-making process. We then investigate the indirect implications of the model for the duration of recommendations and market reactions to recommendations.

Table 4 reports the parameter estimates. Columns 1 to 5 present restricted models for the three-tier rating system and Column 6 presents the full model. Column 7 presents the full model for the five-tier rating system. The bottom two rows show the Brier score and the logarithm of the marginal likelihood of a particular model used to assess the goodness of model fit.

Due to the ordinal nature of stock recommendations, the ordered probit model has been widely used to analyze recommendations. It captures an "idealized" analyst who introduces no recommendation frictions and who has no persistent private information. The ordered probit model is nested in our framework when both δ and ρ are set to zero. This model essentially corresponds to that in Conrad et al. (2006). Columns 1 and $1'$ present parameter estimates of an ordered probit model obtained using our MCMC approach and conventional maximum likelihood, respectively. The two methods yield nearly identical parameter estimates. We defer discussion of the publicly-available determinants of stock valuation to the full model. Figure 1 in the introduction depicts the equilibrium bins for this model.

The ordered probit model fits the data *poorly*. This poor fit is reflected in the large Brier score, which indicates a large discrepancy between predicted probabilities of recommendations and actual outcomes. Moreover, the ordered probit model fitted by maximum likelihood yields a pseudo-R^2 of only 2.09%. The poor performance of the ordered probit model in explaining analyst recommendations is commonly reported in analyses of both recommendation levels and recommendation changes, where the pseudo-R^2 is invariably less than 5% (see Conrad et al., 2006, or Ljungqvist et al., 2007).[10]

It is important to emphasize that our approach is agnostic about what enters an idealized analyst's model of valuation—we let the data tell us how different components enter. The key for the poor fit is that no matter how the model of valuation is formulated, it predicts far too many recommendation revisions. Figure 1 in the introduction presents a sample valuation path without persistent information, and the consequent impacts on recommendations over a 36 month period. This model predicts 20 recommendation revisions over this short period (and there would be many more were one to use weekly, rather than monthly, observations, or to consider five-tier brokerages). In essence, while there is some persistence in recommendations due to persistence in public information data and quarterly arrival of earnings information (i.e., there is persistence in most firm and analyst fundamentals in X), there is *far* too little to generate the infrequent revisions of recommendations found in the data.

Column 2 considers an idealized analyst who does not set recommendation revision frictions, but does gather information that the econometrician does not have, information that has persistent valuation implications. The autoregressive coefficient estimate for this information source is *very* high, $\hat{\rho} = 0.90$, and hugely credible/significant. Incorporating this persistent information source cuts the Brier score almost in half, from 0.34 to 0.18, and it is accompanied by an enormous Bayes factor of exp(49740): there is a *vast* improvement in model fit once one accounts for the temporal correlation in an analyst's information.

The high estimated persistence in an analyst's information reduces the frequency of recommendation revisions, but only to a limited extent. To see why, consider a stock with a valuation in the previous month of $V_{t-1}^* = 4.2$, which is well above the Buy rating cutoff of 3.5. If an analyst receives a one standard deviation positive information shock ($\varepsilon_t = +1$), raising V_t^* to 5.2,[11] then the slow decay in the valuation consequences of this information means that it is likely to take a long time for the analyst's valuation to drop out of the buy bin, even if some negative public information news arrives. Conversely, a one-standard deviation

[10]The analogous regression in which, as in Loh and Stulz (2011), *changes in recommendation* are regressed on the same right-hand variables as those in the ordered probit fits the data even less well, with an adjusted-R^2 of only 0.33%.

[11]For the purpose of this illustration, assume u_{t-1} is 0.

negative shock ($\varepsilon_t = -1$) would lead to a downgrade to Hold as V_t^* drops to 3.2. However, the analyst's valuation would slowly revert, rising due to the decay of ε_t. Hence, absent arrival of other information about the stock's value, the analyst would switch back to a buy recommendation after 4 months. Thus, while the high persistence in an analyst's information reduces the frequency of recommendation revisions, it can only do so to a limited extent.

Figure 3 depicts the equilibrium bins for an idealized analyst who has persistent information, and it illustrates sample valuation paths *for the same common public information valuation path and private information shocks as Figure 1.* Persistent information reduces the number of recommendation revisions from 20 to 12. This reduction helps explain why the model fit is so much better, but also why the data remain badly described by an idealized analyst, even when his information exhibits tremendous persistence. No matter how much persistence is built into the analyst's model of valuation, the idealized analyst model *cannot* deliver low likelihoods of recommendation revisions whenever the current valuation is close to a ratings cutoff: slight changes in prices would lead to frequent recommendation revisions as it repeatedly crosses the cutoff, especially were one to use higher frequency observations.

To generate the substantial stickiness/path dependence in recommendations implicit in the small off-diagonal transition probabilities shown in Table 1, one needs frictions in recommendation revisions strategically introduced by analysts who value intertemporal consistency in recommendations. Columns 3 to 6 present estimates of models in which analysts introduce such frictions. Column 3 presents estimates for a model with (a) only two friction parameters, $\delta_\uparrow$ and $\delta_\downarrow$, one for upward revisions and one for downward revisions, where (b) analysts have no persistent private information. Both revision friction estimates are *large* ($83-87\%$ of the size of the hold bin) and highly significant. Model 3 has a *far* better goodness of fit than model 2 (Bayes factor of $exp(75014)$, and the Brier score falls from 0.18 to 0.15), even though model 3 lacks persistent private information. This emphasizes that to reconcile the dynamics of analyst recommendations, one must incorporate the strategic reluctance of analysts to revise recommendations: revision frictions are even more important drivers of an analyst's recommendation decision processes than persistent private information. Importantly, there would be little impact on estimates were one to use higher frequency (e.g., bi-weekly) recommendation observations, because valuations rarely change sharply over short windows: moderate changes in valuations *cannot* lead to successive changes in recommendations. In this way, recommendation revision frictions capture any temporal stickiness in recommendations.

Table 4: Parameter Estimation

Model	(1)	(1')	(2)	(3)	(4)	(5)	(6)	(7)
Recommendation Cutoffs								
μ_5								4.340
								(526.7)
μ_4	1.559	1.565	3.498	1.782	1.745	1.787	1.966	2.910
	(384.3)	(405.6)	(364.4)	(370.0)	(339.3)	(392.9)	(455.9)	(374.1)
μ_3	0	0	0	0	0	0	0	1.154
								(140.33)
μ_2								0
Persistence of Information								
ρ			0.895		0.401	0.389	0.313	0.609
			(607.4)		(85.47)	(83.73)	(66.15)	(116.9)
Recommendation Revision Frictions								
$\delta_{\uparrow}$				1.445	1.337	1.346		
				(974.0)	(266.9)	(262.7)		
$\delta_{\downarrow}$				1.631	1.525	1.553		
				(954.7)	(403.4)	(394.5)		
$\delta_{5\uparrow}$								1.174
								(43.17)
$\delta_{5\downarrow}$								1.324
								(151.6)
$\delta_{4\uparrow}$							1.336	1.037
							(147.5)	(165.3)
$\delta_{4\downarrow}$							1.563	1.209
							(291.3)	(75.29)
$\delta_{3\uparrow}$							1.787	1.509
							(358.2)	(104.0)
$\delta_{3\downarrow}$							0.822	0.814
							(91.11)	(41.52)
$\delta_{2\uparrow}$								1.136
								(117.1)

Continued on next page

Table 4 – continued from previous page

parameter	(1)	(1')	(2)	(3)	(4)	(5)	(6)	(7)
$\delta_{2\downarrow}$								1.207
								(10.29)
Publicly-available Characteristics								
$const$	0.713	0.713	1.092	0.768	1.480	0.690	1.251	2.897
	(18.02)	(15.76)	(12.43)	(14.82)	(375.84)	(11.48)	(12.61)	(12.53)
ret_{-1}	0.001	0.001	−0.000	−0.001		−0.002	−0.004	−0.002
	(4.63)	(4.58)	(−1.72)	(−4.83)		(−6.93)	(−8.58)	(−4.39)
$ret_{-2:-6}$	0.001	0.001	0.001	0.001		0.001	0.001	0.001
	(4.92)	(5.00)	(4.24)	(6.58)		(5.15)	(3.84)	(1.25)
$ret_{-7:-12}$	0.002	0.002	0.002	0.002		0.002	0.002	0.001
	(18.77)	(18.94)	(13.08)	(13.57)		(11.31)	(7.13)	(2.85)
$\sigma_{-1:-6}$	0.002	0.002	0.004	0.003		0.003	0.002	−0.001
	(8.68)	(8.76)	(6.28)	(7.79)		(6.54)	(3.08)	(−0.75)
$\log(turnover_{-1:-6})$	−0.042	−0.040	−0.076	−0.056		−0.049	−0.042	−0.026
	(−9.06)	(−8.99)	(−7.01)	(−8.98)		(−6.89)	(−3.75)	(−2.23)
$\log(MktCap_{-1:-6})$	−0.053	−0.053	−0.175	−0.058		−0.072	−0.025	−0.041
	(−19.44)	(−19.47)	(−28.12)	(−16.20)		(−17.34)	(−3.62)	(−2.48)
$\log(num\_anal)$	0.115	0.115	0.310	0.110		0.134	0.060	0.117
	(22.26)	(22.40)	(23.71)	(15.88)		(16.27)	(4.84)	(4.50)
$HSize$	−0.080	−0.080	−0.167	−0.067		−0.077	−0.051	−0.120
	(−35.34)	(−35.65)	(−31.70)	(−23.11)		(−22.26)	(−9.40)	(−8.81)
SUE	0.035	0.035	0.051	0.011		0.015	0.017	0.034
	(12.78)	(12.83)	(11.54)	(2.69)		(3.35)	(3.27)	(4.02)
$SUE \times D_{EA}$	0.029	0.030	0.018	0.026		0.024	0.025	0.017
	(6.00)	(6.06)	(4.87)	(5.54)		(5.17)	(4.50)	(2.35)
BM	−0.112	−0.110	−0.256	−0.102		−0.113	−0.080	−0.149
	(−12.61)	(−12.49)	(−13.06)	(−8.50)		(−8.18)	(−3.81)	(−3.40)
EP	0.134	0.135	0.033	0.014		0.006	0.035	0.117
	(3.89)	(3.92)	(0.44)	(0.29)		(0.11)	(0.44)	(0.69)
SG	0.323	0.325	0.679	0.274		0.306	0.177	0.322
	(26.40)	(26.58)	(24.85)	(17.09)		(16.28)	(5.88)	(5.29)

Continued on next page

Table 4 – continued from previous page

parameter	(1)	(1')	(2)	(3)	(4)	(5)	(6)	(7)
ROA	−0.169	−0.167	−0.245	−0.081		−0.116	−0.074	−0.415
	(−5.41)	(−5.39)	(−3.46)	(−1.95)		(−2.43)	(−0.97)	(−2.59)
$FAge$	−0.001	−0.001	−0.003	−0.001		−0.001	−0.001	−0.003
	(−6.04)	(−6.08)	(−9.68)	(−4.99)		(−5.45)	(−0.95)	(−2.57)
$FRtoP$	1.299	1.312	2.330	1.269		1.442	1.434	1.705
	(5.79)	(5.86)	(6.00)	(4.15)		(4.29)	(3.88)	(2.01)
$CFtoP$	2.352	2.347	3.901	2.400		2.469	2.342	2.260
	(10.77)	(10.74)	(9.31)	(8.07)		(7.32)	(4.65)	(2.40)
$FDisp$	−8.411	−8.529	−18.145	−10.867		−10.619	−5.430	−6.975
	(−8.11)	(−8.24)	(−10.07)	(−7.64)		(−6.78)	(−2.42)	(−3.86)
$FDev$	5.493	5.480	13.093	7.188		8.526	8.684	13.091
	(11.44)	(11.38)	(15.52)	(11.04)		(11.72)	(8.13)	(7.46)
IH	0.257	0.256	0.388	0.231		0.221	0.217	0.271
	(21.32)	(21.31)	(13.53)	(14.27)		(11.73)	(7.48)	(4.16)
D_{IB}	0.087	0.087	0.180	0.088		0.095	0.088	0.118
	(6.73)	(6.75)	(6.40)	(5.64)		(5.45)	(4.92)	(6.28)
$\log(year\_brkg)$	−0.066	−0.065	−0.106	−0.056		−0.057	−0.057	−0.054
	(−15.03)	(−14.98)	(−10.57)	(−9.71)		(−8.63)	(−5.11)	(−3.19)
$\log(year\_IBES)$	0.028	0.027	0.047	0.033		0.032	0.043	−0.026
	(5.57)	(5.50)	(3.94)	(4.96)		(4.15)	(3.48)	(−0.99)
Goodness of Fit								
$\#obs$	241076	241076	241076	241076	241076	241076	241076	89726
$\#pair$	8224	8224	8224	8224	8224	8224	8224	3059
$Brier\ score$	0.338		0.180	0.151	0.147	0.144	0.141	0.197
LML	−182194		−132454	−107180	−100534	−93586	−86452	−43881
$pseudo-R^2$		2.09%						

The ratio of the posterior mean to standard deviation is reported in parentheses. The row labeled "LML" shows the logarithm of a model's marginal likelihood, $\log(\Pr(D|M))$.

Figure 3: Idealized analyst with persistent information

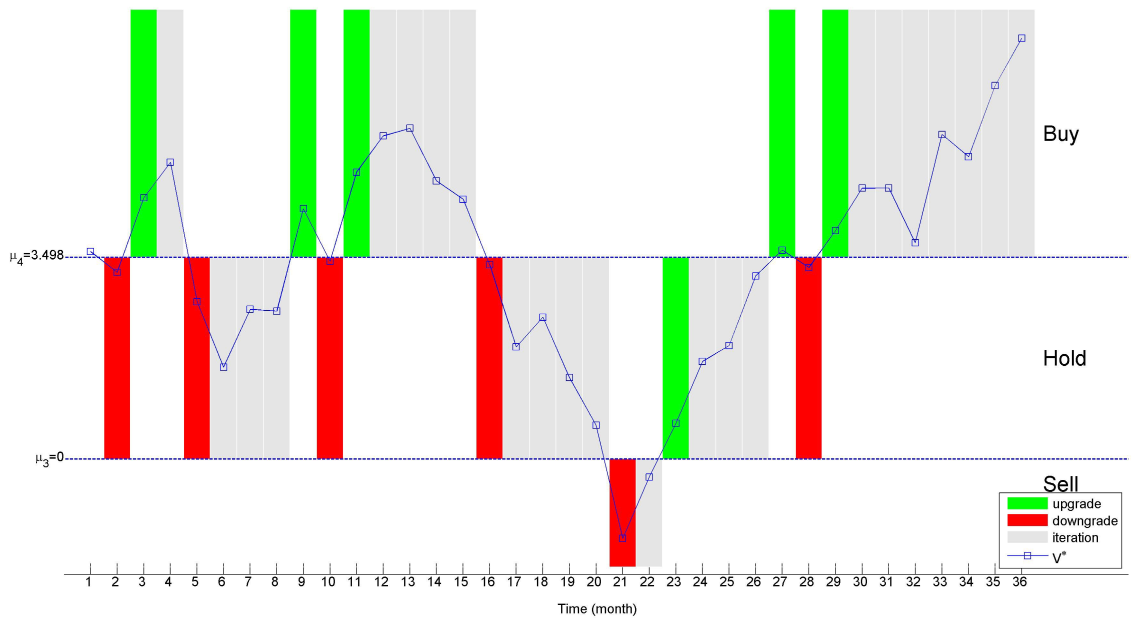

Recommendation cutoffs, and sample valuation and recommendation paths over a 36-month period for an idealized analyst who has persistent private information. The sample valuation path uses *the same common public information valuation path and private information shocks as Figure 1.*

Column 4 presents a model with the same two recommendation frictions, $\delta_\downarrow$ and $\delta_\uparrow$, *and* persistent private information for analysts, where we now *discard all* publicly-available information save for the constant. Despite throwing away *all* public information, there is a massive improvement in model fit from model 3 to model 4 (Bayes factor of $exp(6646)$). This underscores the collective importance of both persistence in analyst information and recommendation frictions for the recommendation process. Both the revision frictions and information persistence parameters are highly significant. These two sources of stickiness in recommendations serve as substitutes in slowing down the frequency of recommendation revisions: the presence of ρ reduces estimates of $\delta_\uparrow$ and $\delta_\downarrow$ by roughly 7%; while the presence of $\delta_\downarrow$ and $\delta_\uparrow$ slashes the estimate of ρ by 55%. Phrased differently, failing to account for both sources of stickiness in recommendations biases estimates significantly.

Column 5 presents estimates of model 5, which features persistent analyst information and recommendation revision frictions, augmented to include the public information variables from model 1. The Bayes factor for models 3 and 5 of $exp(13324)$ is again extremely large: persistent analyst information and recommendation revision frictions are largely complementary sources of improved model fit. These complementarities are also indicated by the huge ratios of the mean to standard deviation of the parameter estimate: 85.5 for the information persistence parameter ρ, 266.9 for $\delta_{\uparrow}$ and 403.4 for $\delta_{\downarrow}$. Phrased differently, persistent private information and recommendation revision frictions capture fundamentally different economic phenomena—*persistence in analyst information (or in the analyst's assessment of value) is not a proxy for an unwillingness of analysts to revise recommendations*—and both are central to understanding the dynamics of analyst recommendations.

The improvement in the goodness of fit relative to model 4, which has no public information, is large (Bayes factor of $exp(6948)$), but it pales in comparison with the improvements associated with introducing revision frictions or persistent analyst information. Moreover, the predictive improvement suggested by the 0.003 reduction in the Brier score is tiny. In sum, standard public information sources matter far less for the dynamics of analyst recommendations than do revision frictions and persistent analyst information.

Column 6 presents estimates for our full model, in which recommendation revision frictions are cutoff-specific and analysts have persistent private information. Recommendation revision frictions $\delta_{k\downarrow}$ and $\delta_{k\uparrow}$ bear the same index k as cutoff μ_k: $\delta_{k\downarrow}$ is the friction for downgrades from k to $k-1$, while $\delta_{k\uparrow}$ is the friction for upgrades from $k-1$ to k. Analysts do not need to use symmetric recommendation revision frictions to reduce the frequency of recommendation revisions; they can tailor them to reflect other considerations (see Figure 4).

The full model fits the recommendation data the best. It has the smallest Brier score of 0.141 and the large Bayes factor of $exp(7134)$ versus model 5 provides conclusive evidence against the other models. Figure 5 depicts the equilibrium bins for the full model, and it illustrates sample valuation paths *for the same common public information valuation path and private information shocks as Figures 1 and 3.* The figure hints at why the model fit is so much better: only 4 revised recommendations are issued over the 36 month period. Inspection of the estimates of the cutoff-specific revision frictions provides more insights: their magnitudes vary sharply, indicating that the more restricted models are significantly mis-specified. In particular, frictions out of hold are much smaller than those into hold, and a single one-directional friction cannot simultaneously account for this. An indirect indicator is that the hold bin size increases by 12%, so that the frictions are smaller relative to the bin size, and the more flexible formulation of the recommendation-specific revision frictions reduces the analyst's information as a source of persistence in recommendations by almost 25%.

Column 7 presents estimates for brokerages that use the traditional five-tier rating sys-

Figure 4: Cutoff and recommendation revision friction estimates

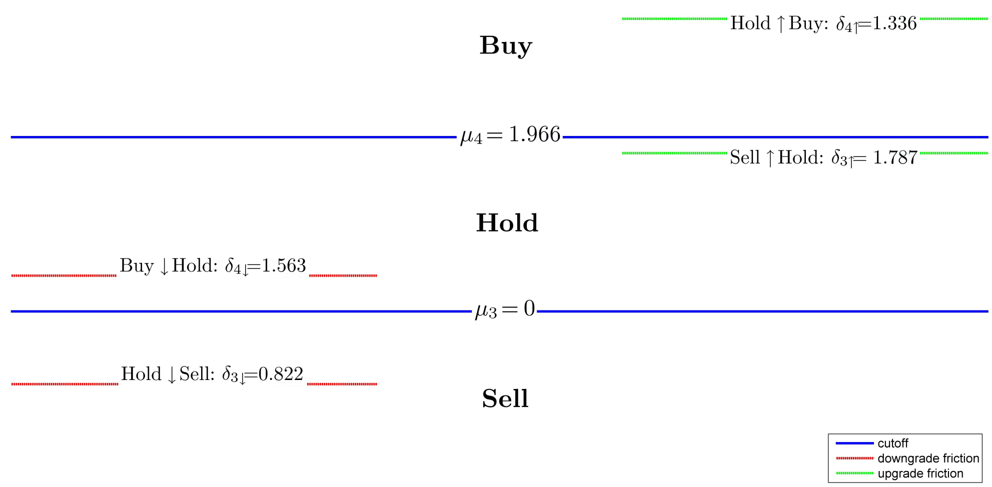

Cutoff and recommendation revision friction estimates of the full model, in which revision frictions are cutoff-specific, and analysts have persistent information. Cutoff-specific revision frictions, $\delta_{k,\downarrow}$ and $\delta_{k\uparrow}$, bear the same index k as cutoff μ_k: $\delta_{k,\downarrow}$ is the friction for downgrades *from* k to $k-1$, while $\delta_{k\uparrow}$ is the friction for upgrades from $k-1$ *to* k.

tem. Most estimates are qualitatively similar to those for the three-tier system. For example, the recommendation revision friction from sell to hold is far higher than that from hold to sell. Observe, too that the frictions from strong sell to sell and strong buy to buy are very large relative to the sizes of the sell and buy bins (98% and 92% respectively), so that most revisions from strong buy and strong sell are to hold. Thus, the overall nature of the recommendation revision frictions that five-tier brokerages introduce qualitatively lead them in the direction of behaving like three-tier brokerages.

This does *not* imply that brokerages using a five-tier rating system behave like their three-tier counterparts. There are at least two reasons why one cannot pool strong buys with buys, strong sells with sells, and estimate a homogeneous brokerage using a three-tier rating system. First, the estimates show why brokerage houses using a five-tier rating system tend to have more optimistic assessments of the firms that they follow: the average new recommendation initiation for those using five ratings is roughly at the buy/hold cutoff (the constant is 2.9), while that for those using only three ratings is at the 62^{nd} percentile of the hold recommendation bin above the sell cutoff (the constant is 1.3). This may reflect that

Figure 5: Pretty good analyst

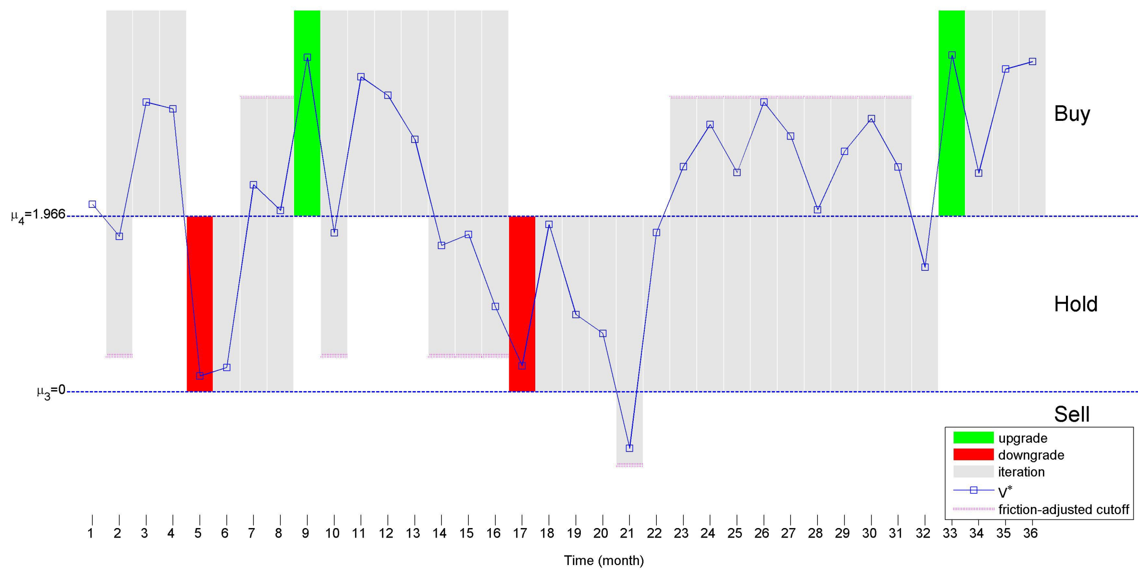

Recommendation cutoffs, and sample valuation and recommendation paths over a 36-month period for a pretty good analyst who has persistent information. The sample valuation path uses the *same common public information valuation path and private information shocks as Figures 1 and 3.*

they have very different audiences: five-tier brokerages (which tend to be smaller, without an investment bank side) may be oriented more toward appealing to retail investors. Second, the estimate of information persistence is much higher for the five-tier brokerage houses.[12]

Our estimates highlight how analysts design bins and revision frictions to generate trades from satisfied retail investors:

- The large frictions out of strong buy and strong sell plausibly reflect that revisions to hold generate trading activity by inducing investors to unwind positions, but revisions from strong buy to buy that maintain a positive assessment, or from strong sell to sell that maintain a negative assessment do not.

- Revision frictions *out* of hold are small, suggest that *maintained* hold ratings (as op-

[12] One might worry that this high estimate could indicate that some analysts in the five-tier brokerage sample behave as if they were at three-tier brokerages. Such mis-classification would lead to higher estimates of revision frictions within buy and sell categories, and greater information persistence. However, the many transitions from buy to strong buy and strong buy to buy suggest that the extent of any misclassification is slight.

posed to *revisions* to hold) discourage investors from trading. Analysts can limit repeated revisions both with symmetric and asymmetric frictions. Analysts prefer to design frictions in and out of hold asymmetrically, with smaller frictions out of hold, so that recommendations spend "less time" in hold, thereby generating more client trades.

Even though hold revision frictions are small, the model delivers the prevalence of hold recommendations (39% for five-tier brokerages, 49% for three-tier brokerages) in three key ways:

- The hold recommendation bin for five-tier brokerages is large, about 25% larger than the buy bin, and 50% larger than the sell bin.
- The estimated average firm for five-tier brokerages is roughly at the buy-hold recommendation cutoff, which places extensive probability mass on an initial hold recommendation; and the estimated average firm for three-tier brokerages is slightly above the hold bin midpoint.
- Analysts at five-tier brokerages are less likely to face revision frictions from buy or sell into hold due to the high frictions from strong sell to sell and strong buy to buy, which results in most transitions going from strong buy and strong sell straight to hold.[13]

These findings make economic sense on other fronts, as well. That the average firm for which coverage is initiated is a hold, but closer to a buy than a sell, supports the notion that analysts tend to follow stocks that they deem to have better prospects. This is consistent with their retail clients being less likely to short-sell. As a result, covering firms with poorer prospects generates fewer retail client orders from which brokerage houses can profit. At the same time, analysts want clients to profit from trades—a happy client is likely to trade again—so analysts want there to be meaning to buy and sell recommendations, and hence are reluctant to issue such recommendations unless profits are somewhat likely to result.

We now turn to the publicly-available determinants of value. Of note, we see that in contrast to existing findings, once we control for a pretty good analyst's information and revision frictions, better lagged return measures of past firm performance cease to systematically raise the analyst's assessment: the one month lagged return now enters negatively,

[13] These estimates also give insight into where improvement in model fit occurs versus more restricted models of five-tier brokerages. Recall that strong sells only comprise 1.4% of the sample, so that the two frictions, $\delta_{\uparrow}$ and $\delta_{\downarrow}$ do not weigh transitions from strong sells heavily in the estimation. As a result, $\delta_{\uparrow}$ is far larger than $\delta_{2\uparrow}$. In turn, the large size of $\delta_{\uparrow}$ necessitates a large value for μ_2 (so that transitions from strong sell to sell can occur with positive probability); but then the model cannot deliver few initial sell recommendations.

while more distant returns enter slightly positively. Qualitatively, Column 6 reveals that a pretty good analyst has higher assessments of firms:

- for which the analyst has relatively higher estimated forecasts of earnings (vis à vis the consensus), consistent with Womack (1996) or Jegadeesh et al. (2004).
- that have positive earnings surprise, especially in the earnings announcement window in which they are reported, as in Chan, Jegadeesh and Lakonishok (1996), Jegadeesh, Kim, Krische and Lee (2004), or Ivkovic and Jegadeesh (2004).
- that draw more attention from other analysts, consistent with more analysts following stocks they believe are undervalued, or possibly analysts valuing a stock's "glamor".
- that are smaller, as measured by higher sales growth (see Lakonishok, Shleifer, and Vishny (1994)) or lower book-to-price ratios (see Jegadeesh et al. 2004)).
- with less turnover, consistent with Lee and Swaminathan (2000), who argue that turnover is a contrarian sign, associated with lower returns.
- about which there is less uncertainty, as captured by forecast dispersion in earnings (see Diether et al. (2002), Zhang (2006)), lesser dispersion in recommendations, or more analyst following or higher institutional holdings.
- for which an analyst's brokerage house has investment banking relationships, consistent with Lin and McNichols (1998), Ljungqvist et al. (2007), O'Brien, McNichols and Lin (2005), Jackson (2005), Cowen, Groysberg, and Healy (2006) and Lim (2001).
- if an analyst is at a smaller brokerage house. Analysts at smaller brokerages also issue more optimistic earnings forecasts (Bernhardt et al. (2006)), and follow smaller firms.
- if an analyst is new at his or her brokerage firm. The coefficient on an analyst's duration at a brokerage firm is negative and highly significant. This is consistent with analysts initially issuing optimistic recommendations in order to generate trading activity, while senior analysts issue more conservative recommendations to preserve reputations.

Quite generally, accounting for revision frictions and information persistence in analyst decision-making sharply reduces the statistical significance/credibility of parameter estimates (relative to the ordered probit model of an idealized analyst), typically by factors of two to five, and the magnitudes of parameter estimates tend to be reduced, as well.

Robustness of Homogeneous Model of Recommendations. Our econometric model of how analysts form recommendations presumes that sources of heterogeneity between analysts or between the firms for which analysts issue recommendations only enter via the assessment of value, V. To assess the validity of this premise, we estimate separate models for subsamples of analysts and firms where one might suspect that analysts' recommendations might vary—over time, by brokerage house size, analyst following and analyst experience. These robustness tests reveal *remarkable* consistency in our estimates.[14] Column 1 of Table 5 reproduces estimates from the full sample. Subsequent columns present estimates for the subsamples of (S1) the second half of the sample period (2006-2010); (S2) large brokerages that on average employed at least 52 analysts over the sample period; (S3) heavily-followed stocks that were covered on average by at least 15 analysts over the sample period; and (S4) senior analysts who have been employed by the same brokerage firm for at least five years. The sample criteria were chosen so that the subsample would have roughly half of the original observations.

We see true intertemporal consistency—comparing columns (Full) and (S1) reveals almost no variation in estimates—there is **no** evidence that analysts have altered how they issue recommendations over this period. Similarly, the subsample of larger brokerage houses (S2), and senior analysts (S4) have very similar estimates, with slightly higher estimates (less than 10% more) of the persistence in an analyst's information, and somewhat higher frictions for recommendation revisions from hold to sell. Analyst's information appears to be slightly more persistent yet for heavily-followed stocks, but even this difference is less than 25%, and the other structural recommendation parameters differ by far less. In sum, the homogeneous model of recommendation formation and revision describes analyst behavior extremely well.

Delayed Incorporation of Information by Analysts? Raedy, Shane and Yang (2006) uncover evidence suggesting that it takes time for analysts to process new information, which would lead them to under-react to it. This leads us to modify our model to estimate the extent to which analysts fully process new information, providing a direct estimate of the amount by which analysts under-react to new information. A by-product of this analysis is that it assuages concerns that such delayed incorporation of information can underlie the high estimates of persistence in analyst information or the large recommendation revision frictions that we find. To do this, we estimate a model in which, of the new private information ε_{ijt}

[14] The very large ratios of the posterior mean to standard deviation that we find for our structural parameters indicate that they are extremely precisely estimated. As a check on this precision, in unreported results, we estimate the model for the subsample of stocks with odd CUSIP numbers. All structural parameter estimates differed by less than 0.01.

that analyst i receives about stock j at time t, the analyst only incorporates a fraction ζ. As a result, the valuation consequences of the analyst's persistent information evolve according to

$$u_{ijt} = \rho(u_{ij,t-1} + (1-\zeta)\varepsilon_{ij,t-1}) + \zeta\varepsilon_{ijt}.$$

The last column of Table 5 presents estimation results for the model in which analysts under-react to new information. This model displays a somewhat improved fit. The estimates indicate that analysts incorporate the vast bulk of new information immediately, incorporating all but 9 percent of information when it arrives. Importantly, the estimates reveal that allowing for delayed incorporation of information does not lower the estimate of information persistence; rather, it raises the estimate by about one third. Moreover, while three of the four estimates of the revision frictions are reduced, all changes in estimates are modest, on the order of 5 to 15%. That the revision frictions remain large was to be expected, as delayed incorporation of information alone could not prevent repeated revision when valuation-price ratios are close to cutoffs. Thus, the estimates indicate that analysts are "close to rational" in their assessments of new information, and that our qualitative findings are largely reinforced by integrating this source of modest "irrationality".

Lost Interest? Columns (6m), (9m) and (15m) present estimates when we use alternative cutoffs of 6, 9 or 15 months for the time after which we assume that an analyst has ceased following a stock absent a recommendation revision or reiteration. Longer windows include more analysts who have ceased following a stock, and hence spuriously suggest stickiness, while shorter windows exclude more analysts who are following a stock, but have not reiterated, spuriously suggesting too little stickiness. We see that any reasonable cutoff level has only modest effects on estimates of structural parameters—as expected, longer windows for continued coverage slightly raise estimates of persistence and revision frictions.

Indirect Evidence. The goodness of fit measures provide conclusive evidence that our model of a pretty good analyst does a vastly superior job of explaining the dynamics of analyst recommendations than do conventional discrete choice models. We now derive indirect implications of our model, and document the associated confirming empirical evidence.

Table 5: Subsample Analysis

Model	(Full)	(S1)	(S2)	(S3)	(S4)	(6m)	(9m)	(15m)	(delay)
Recommendation Cutoffs									
μ_4	1.966	1.957	1.914	1.888	1.918	1.987	1.978	1.958	1.732
	(455.9)	(512.86)	(238.5)	(369.5)	(417.9)	(507.7)	(489.7)	(509.0)	(293.3)
μ_3	0	0	0	0	0	0	0	0	0
Persistence of Information									
ρ	0.313	0.309	0.342	0.388	0.332	0.289	0.300	0.321	0.430
	(66.15)	(47.96)	(48.30)	(51.93)	(43.85)	(48.71)	(59.83)	(75.75)	(89.66)
delayed incorporation of information									
ζ									0.907
									(336.7)
Recommendation Revision Frictions									
$\delta_{4\uparrow}$	1.336	1.342	1.362	1.400	1.320	1.075	1.220	1.395	1.209
	(147.5)	(112.84)	(104.4)	(69.15)	(67.89)	(100.2)	(185.1)	(332.8)	(220.4)
$\delta_{4\downarrow}$	1.563	1.569	1.562	1.509	1.553	1.536	1.561	1.589	1.335
	(291.3)	(263.02)	(165.7)	(126.1)	(138.1)	(272.3)	(263.5)	(239.9)	(256.1)
$\delta_{3\uparrow}$	1.787	1.756	1.757	1.695	1.746	1.747	1.775	1.820	1.535
	(358.2)	(238.41)	(180.9)	(166.3)	(174.1)	(194.4)	(324.3)	(327.6)	(375.3)
$\delta_{3\downarrow}$	0.822	0.836	0.965	0.921	0.889	0.691	0.786	0.848	0.851
	(91.11)	(48.52)	(51.60)	(30.33)	(36.65)	(51.56)	(89.69)	(83.81)	(221.6)
Publicly-available Characteristics									
ret_{-1}	−0.004	−0.003	−0.004	−0.002	−0.002	−0.004	−0.004	−0.004	−0.004
	(−8.58)	(−5.25)	(−6.28)	(−3.66)	(−3.12)	(−8.43)	(−9.18)	(−9.15)	(−9.75)
$ret_{-2:-6}$	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
	(3.84)	(3.77)	(2.61)	(2.08)	(2.76)	(4.03)	(4.32)	(3.35)	(2.72)
$ret_{-7:-12}$	0.002	0.002	0.001	0.002	0.002	0.001	0.001	0.001	0.001
	(7.13)	(6.27)	(4.20)	(3.91)	(4.78)	(5.16)	(6.33)	(6.56)	(6.57)
$\sigma_{-1:-6}$	0.002	0.001	0.002	0.002	0.003	0.000	0.001	0.002	0.002
	(3.08)	(0.80)	(1.54)	(1.76)	(2.37)	(0.02)	(1.71)	(2.66)	(2.68)
$\log(turnover_{-1:-6})$	−0.042	−0.053	−0.013	−0.007	−0.064	−0.034	−0.037	−0.039	−0.035
	(−3.75)	(−3.41)	(−0.77)	(−0.29)	(−3.51)	(−2.31)	(−3.02)	(−3.71)	(−2.95)

Continued on next page

Table 5 – continued from previous page

parameter	(6)	(S1)	(S2)	(S3)	(S4)	(6m)	(9m)	(15m)	(delay)
$\log(MktCap_{-1:-6})$	−0.025	−0.007	−0.081	−0.080	−0.025	−0.010	−0.019	−0.029	−0.026
	(−3.62)	(−0.83)	(−7.83)	(−6.59)	(−2.32)	(−1.19)	(−2.60)	(−4.47)	(−3.47)
$\log(num\_anal)$	0.060	0.043	0.099	0.057	−0.020	0.040	0.054	0.057	0.059
	(4.84)	(2.56)	(5.14)	(1.60)	(−1.01)	(2.47)	(4.01)	(4.82)	(4.47)
$HSize$	−0.051	−0.033	−0.370	−0.045	−0.036	−0.015	−0.036	−0.061	−0.051
	(−9.40)	(−4.71)	(−14.04)	(−4.48)	(−4.16)	(−2.37)	(−6.46)	(−11.59)	(−8.61)
SUE	0.017	0.005	0.014	0.014	0.019	0.017	0.015	0.015	0.018
	(3.27)	(0.70)	(2.83)	(2.42)	(2.27)	(2.76)	(2.85)	(3.08)	(3.69)
$SUE \times D_{EA}$	0.025	0.029	0.010	0.020	0.016	0.018	0.027	0.022	0.024
	(4.50)	(3.72)	(2.23)	(2.22)	(1.74)	(2.46)	(4.38)	(4.01)	(4.69)
BM	−0.080	−0.090	−0.102	0.008	−0.145	−0.044	−0.062	−0.071	−0.083
	(−3.81)	(−3.50)	(−3.15)	(0.19)	(−4.14)	(−1.60)	(−2.70)	(−4.13)	(−3.70)
EP	0.035	0.052	0.012	0.071	0.297	−0.057	0.019	0.013	0.034
	(0.44)	(0.55)	(0.09)	(0.45)	(2.10)	(−0.52)	(0.21)	(0.16)	(0.41)
SG	0.177	0.218	0.173	0.169	0.209	0.186	0.158	0.209	0.173
	(5.88)	(5.30)	(3.79)	(3.09)	(4.13)	(4.83)	(5.02)	(7.45)	(5.42)
ROA	−0.074	−0.199	−0.105	−0.061	0.150	0.038	−0.092	−0.032	−0.089
	(−0.97)	(−1.95)	(−0.87)	(−0.43)	(1.10)	(0.38)	(−1.11)	(−0.44)	(−1.11)
$FAge$	−0.001	−0.000	−0.001	−0.001	−0.001	0.000	−0.001	−0.001	−0.001
	(−0.95)	(−0.84)	(−1.28)	(−1.29)	(−0.21)	(0.60)	(−1.21)	(−2.49)	(−0.92)
$FRtoP$	1.434	1.638	0.871	−0.473	1.435	1.453	0.905	0.452	0.372
	(3.88)	(3.09)	(1.18)	(−0.51)	(3.54)	(2.26)	(1.71)	(0.97)	(0.77)
$CFtoP$	2.342	2.041	3.837	1.710	4.365	0.800	1.776	2.264	2.294
	(4.65)	(3.45)	(5.02)	(1.81)	(5.13)	(1.16)	(3.24)	(4.88)	(4.56)
$FDisp$	−5.430	−5.641	−0.646	−15.832	−5.894	−4.856	−7.083	−2.732	−4.554
	(−2.42)	(−2.13)	(−0.19)	(−3.56)	(−1.56)	(−1.65)	(−2.86)	(−1.30)	(−2.09)
$FDev$	8.684	7.557	7.422	8.198	9.134	6.909	7.163	8.911	9.240
	(8.13)	(5.87)	(4.68)	(4.12)	(5.28)	(4.94)	(6.18)	(8.89)	(8.62)
IH	0.217	0.228	0.284	0.151	0.157	0.145	0.194	0.227	0.203
	(7.48)	(5.95)	(6.40)	(2.68)	(3.33)	(3.79)	(6.19)	(8.30)	(6.67)
D_{IB}	0.088	0.083	0.033	0.110	0.046	0.045	0.064	0.077	0.092

Continued on next page

Table 5 – continued from previous page

parameter	(6)	(S1)	(S2)	(S3)	(S4)	(6m)	(9m)	(15m)	(delay)
	(4.92)	(3.15)	(2.50)	(1.35)	(2.81)	(1.00)	(1.46)	(1.66)	(1.86)
$\log(year\_brkg)$	−0.057	−0.060	−0.030	−0.040	−0.014	−0.065	−0.055	−0.050	−0.056
	(−5.11)	(−4.22)	(−1.64)	(−2.12)	(−0.50)	(−4.87)	(−4.21)	(−4.79)	(−4.73)
$\log(year\_IBES)$	0.043	0.062	−0.004	0.067	0.085	0.054	0.039	0.018	0.045
	(3.48)	(3.79)	(−0.19)	(3.01)	(2.72)	(3.41)	(3.07)	(1.59)	(3.41)
const	1.251	1.388	1.422	1.345	1.119	1.256	1.301	1.259	1.098
	(12.61)	(10.48)	(10.91)	(3.75)	(4.80)	(10.07)	(12.24)	(13.52)	(10.24)
Goodness of Fit									
#obs	241076	131067	122133	115739	96522	111455	183870	296164	241076
#pair	8224	6541	4426	4049	3119	4002	6346	9914	8224
Brier score	0.180	0.176	0.188	0.201	0.188	0.186	0.180	0.180	0.136
LML	−86452.0	−54345.6	−51174.1	−54328.3	−40946.2	−48341.3	−77329.5	−127554.1	−75618.0

The ratio of the posterior mean to standard deviation is reported in parentheses. The row labeled "*LML*" shows the logarithm of a model's marginal likelihood, $\log(\Pr(D|M))$.

We first show how the pretty good analyst model provides a theoretical framework that can reconcile the differential impacts of recommendation revisions made around significant public information events, such as quarterly earnings announcements. As the most important source of firm-specific information, earnings announcements convey material and "lumpy" information content about earnings, and other key firm characteristics (e.g., sales, margins and investment (Brandt et al., (2008)). The earnings guidance that firms provide is another channel for lumpy firm-specific information release. Earnings announcements and guidance often feature large surprises that discontinuously shift analysts' assessments of a firm's value. In contrast, information arrival about firm value at other times tends to be smooth.

This difference in information arrival has implications within our model for the differential impacts of recommendation revisions made inside and outside earnings announcement and guidance windows. Recommendation revisions made outside EA or EG windows typically occur when valuation assessments smoothly cross friction-adjusted recommendation cutoffs. In contrast, revisions issued inside these windows are more likely to be driven by "jumps" in valuations past friction-adjusted recommendation cutoffs due to "lumpy" information release.

It follows that if, as in our model, given an outstanding recommendation, analysts issue higher recommendations to firms they value more highly on a per dollar basis, then revisions issued within EA or EG windows should tend to have longer retracement durations than those issued outside these windows. That is, inside an EA or EG window, the valuation is likely further from the friction-adjusted recommendation cutoff, so it will take longer for an analyst's assessment of value to retrace toward the cutoff for a revision back to the original recommendation. Moreover, the valuation information associated with a recommendation *revision* inside an EA or EG window should greater on average than the valuation information conveyed by a revision made outside of those windows. That is, the CAR (cumulative abnormal return) impact of a recommendation revision should be greater inside a window.

Importantly, this effect should *not* exist for *new* recommendations as no information is conveyed by an earnings announcement about the *location* of an analyst's assessment of value relative to cutoffs. Thus, via this difference in differences, we can control for CAR impacts of information arrival in earnings announcements or guidance that are not due to recommendation revisions (see Ivkovic and Jegadeesh (2004) and Kecské et al. (2010)).

We focus on recommendation revisions of one-level, as revisions of multiple levels (e.g., from a buy to a sell) necessarily reflect discontinuities in the valuation consequences of new information, so we cannot conclude that valuations tend to be closer to cutoffs following recommendations outside EA and EG windows. For revisions of multiple levels, our model only predicts that they are more likely within these windows than outside, which we find in the data. The EA window is defined as the three-day period on and after the date of a firm's quarterly earnings announcement, where announcements after close are treated as if they occurred on the next trading day. We then determine whether the date of a recommendation[15] is inside or out of an EA window. We obtain earnings guidance dates from the First Call Guidelines database and define three-day earnings guidance windows in the analogous way.

Table 6: Retracement durations of revisions inside vs. outside EA windows

Panel A. Retracement durations of one-level revisions

EA Window	Obs.	Mean	Std. Dev.	95% Conf. Interval
In	6154	158.268	96.924	(155.845, 160.689)
Out	19402	152.064	98.064	(150.683, 153.443)
diff.		6.204*** (4.34)		

Panel B. Retracement durations of one-level upgrades

EA Window	Obs.	Mean	Std. Dev.	95% Conf. Interval
In	3021	163.396	96.563	(159.951, 166.841)
Out	9296	159.705	98.659	(157.699, 161.711)
diff.		3.691** (1.80)		

Panel C. Retracement durations of one-level downgrades

EA Window	Obs.	Mean	Std. Dev.	95% Conf. Interval
In	3133	153.322	97.028	(149.924, 156.721)
Out	10106	145.034	96.987	(143.143, 146.926)
diff.		8.288*** (4.18)		

The row "diff." refers to the difference in average retracement durations of revisions made inside vs. outside EA windows. t-statistics of tests on the equality of means are reported in parentheses. *** and ** denote statistical significance at the 1% and 5% levels.

Panel A of Tables 6 and 7 reports summary statistics of durations (in days) for a recommendation revision to retrace to its original level, both for revisions made inside the associated windows and for revisions outside of *both* windows. Consistent with our predictions, on average it takes over 6 days longer for revisions issued inside EA windows to return to their original levels than it does for revisions made outside both windows; and it takes

[15] If the date is not a trading day, we use the next trading date.

Table 7: Retracement durations of revisions inside vs. outside EG windows

Panel A. Retracement durations of one-level revisions				
EG Window	Obs.	Mean	Std. Dev.	95% Conf. Interval
In	3245	159.859	95.671	(156.566, 163.152)
Out	19402	152.064	98.064	(150.684, 153.443)
diff.		7.795*** (4.21)		

Panel B. Retracement durations of one-level upgrades				
EG Window	Obs.	Mean	Std. Dev.	95% Conf. Interval
In	1522	162.179	96.358	(157.335, 167.024)
Out	9296	159.705	98.659	(157.699, 161.711)
diff.		2.474 (0.91)		

Panel C. Retracement durations of one-level downgrades				
EG Window	Obs.	Mean	Std. Dev.	95% Conf. Interval
In	1723	157.809	95.040	(153.318, 162.300)
Out	10106	145.034	96.987	(143.143, 146.926)
diff.		12.775*** (5.07)		

The row "diff." refers to the difference in average retracement durations of revisions made inside vs. outside earnings guidance (EG) windows. t-statistics of tests on the equality of means are reported in parentheses. *** denotes statistical significance at the 1% level.

over 8 days longer for revisions issued inside EG windows to return. These highly significant differences in average lifetimes are roughly 5% of the average duration of a revision before its retracement. Panels B and C decompose revisions into upgrades and downgrades.[16] They show that the effect is stronger for downgrades: it takes almost 13 days longer for downgrades issued inside EG windows to return than for downgrades issued at other times, whereas the difference is only 2 days for upgrades. This likely reflects that negative earnings guidance tends to be larger in magnitude than positive guidance, which may take the form of a firm confirming that earnings should be in line with past guidance.[17]

We next explore the implications for market responses to recommendation revisions inside vs. outside earnings announcement and guidance windows. The market reaction is measured

[16]No systematic differences in retracement durations emerge for three-tier vs. five-tier brokerages.

[17]One can also use retracement durations as a ball park test of our premise that analysts set the same cutoffs μ_j for new and revised recommendations. Retracement durations are only marginally longer for new recommendations than revised ones.

by the three-day CAR following a revision issued by analyst i for stock j at day d,

$$CAR_{ijd} = \prod_{d=0}^{2} R_{jd} - \prod_{d=0}^{2} R_M, \tag{7}$$

where R_{jd} and R_M are the raw stock and market daily return, respectively. Day 0 ($d = 0$) is the I/B/E/S reported recommendation date or the following trading day if the recommendation date is not a trading date. We exclude recommendation revisions made on the same day as an announcement or guidance (or the next day if the EA or EG is after close) to avoid having the CAR reflect both the announcement and the revision.

Table 8: Market reaction to revisions made inside vs. outside EA windows

Panel A. Market reaction to one-level upgrades				
EA Window	Obs.	Mean	Std. Dev.	95% Conf. Interval
In	8221	2.917	7.417	(2.756, 3.077)
Out	24488	2.345	7.915	(2.245, 2.443)
diff.		0.573^{***} (5.95)		

Panel B. Market reaction to one-level downgrades				
EA Window	Obs.	Mean	Std. Dev.	95% Conf. Interval
In	8518	−3.347	8.841	(−3.534, −3.158)
Out	24793	−1.821	8.517	(−1.926, −1.714)
diff.		-1.526^{***} (13.87)		

Three-day cumulative abnormal returns (CAR) associated with recommendation revisions issued inside and outside earnings announcement (EA) windows. The row "diff." refers to the difference in average CAR of revisions made inside vs. outside EA windows. t-statistics of tests on the equality of means are reported in parentheses. *** denotes statistical significance at 1% level.

Tables 8 and 9 report findings for three-day CARs following revisions. Consistent with our predictions, upgrades issued inside EA and EG windows have far larger market impacts than those issued outside of both windows, and these CAR differences are highly significant. Market reactions to downgrades issued inside EA windows have CARs that are 1.5 percentage points lower than the CARs for downgrades issued outside these two windows; and those for downgrades issued inside EG windows are even larger, roughly double those for earnings announcements. CAR differences for upgrades are smaller, roughly one-third of those for downgrades. This likely reflects that bad announcements or guidance conveys "more" news

Table 9: Market reaction to revisions made inside vs. outside EG windows

Panel A. Market reaction to one-level upgrades				
EG Window	Obs.	Mean	Std. Dev.	95% Conf. Interval
In	4163	3.401	7.629	$(3.169, 3.633)$
Out	24488	2.345	7.915	$(2.245, 2.443)$
diff.		1.056*** (7.66)		

Panel B. Market reaction to one-level downgrades				
EG Window	Obs.	Mean	Std. Dev.	95% Conf. Interval
In	4756	-4.802	9.009	$(-5.057, -4.546)$
Out	24793	-1.821	8.517	$(-1.926, -1.714)$
diff.		-2.980*** (20.71)		

Three-day cumulative abnormal returns (CAR) associated with recommendation revisions issued inside and outside earnings guidance (EG) windows. The row "diff." refers to the difference in average CAR of revisions made inside vs. outside EG windows. t-statistics of tests on the equality of means are reported in parentheses. *** denotes statistical significance at 1% level.

than good announcements or guidance. As a result, the market's forecast of the change in an analyst's assessment associated with a downward revision is greater.

One might worry that discarding recommendation revisions made on the same days as earnings announcements or guidance is not enough to isolate the effect of a recommendation revision from that of the announcement or guidance. To address this, we observe that the discontinuity in valuation assessment in EA and EG windows *only occurs for revisions and not for new recommendations*. That is, a *new* recommendation made within an EA or EG window conveys no information about value relative to cutoffs that is distinct from a new recommendation made outside the window. Thus, differences in CARs for new recommendations inside vs. outside these windows control for the direct information arrival associated with earnings announcements or guidance. This leads us to employ a difference-in-difference analysis of revisions vs. new recommendations inside and out of EA and EG windows.

Tables 10 and 11 present the "difference-in-difference" analysis for revisions made inside and outside EA and EG windows. They also provide a more detailed decomposition for different recommendation levels. These tables provide further confirmatory evidence for the prediction. In particular, except for downgrades to sell, all difference-in-differences in CARs are highly statistically significant with the "correct" signs, and their large magnitudes range

Table 10: Difference-in-Difference Analysis of CARS for Earnings Announcements

Upgrade to Buy/Strong Buy

	In EA Win	Out EA & EG Win	diff
Up	3.051***	2.385***	0.666***
	(7.82)	(25.79)	(4.77)
Init	1.487***	1.247***	0.240**
	(4.00)	(37.08)	(2.57)
diff	1.564***	1.137***	0.426***
	(4.70)	(19.86)	(3.22)

Upgrade to Hold

	In EA Win	Out EA & EG Win	diff
Up	2.182***	1.662***	0.520***
	(11.93)	(25.50)	(8.74)
Init	−1.768***	−0.577***	−1.190***
	(14.80)	(15.76)	(13.03)
diff	3.950***	2.240***	1.710***
	(12.24)	(23.99)	(8.77)

Downgrade to Hold

Down	−3.267***	−1.568***	−1.699***
	(8.29)	(18.71)	(5.95)
Init	−1.768***	−0.577***	−1.190***
	(13.77)	(14.62)	(12.09)
diff	−1.499***	−0.991***	−0.509***
	(4.94)	(14.69)	(3.50)

Downgrade to Sell/Strong Sell

Down	−3.755***	−2.430***	−1.325**
	(4.34)	(8.31)	(1.83)
Init	−2.912***	−1.737***	−1.175***
	(6.66)	(17.18)	(4.54)
diff	−0.843***	−0.693***	−0.150
	(3.17)	(4.74)	(0.43)

Difference-in-difference results of CARs for revisions made inside vs. outside earnings announcement windows. The CAR associated with initiations of the corresponding rating is used as the control group. ***, ** and * denote statistical significance at 1%, 5% and 10% levels.

Table 11: Difference-in-Difference Analysis of CARs after Earnings Guidance

Upgrade to Buy/Strong Buy

	In EG Win	Out EA & EG Win	diff
Up	3.574***	2.385***	1.189***
	(9.00)	(26.06)	(6.18)
Init	1.730***	1.247***	0.483***
	(5.21)	(37.47)	(3.83)
diff	1.844***	1.138***	0.706***
	(5.35)	(20.08)	(3.99)

Upgrade to Hold

	In EG Win	Out EA & EG Win	diff
Up	2.230***	1.670***	0.560***
	(11.82)	(25.77)	(9.31)
Init	−2.718***	−0.579***	−2.139***
	(19.32)	(15.91)	(17.86)
diff	4.948***	2.248***	2.700***
	(13.15)	(24.24)	(10.21)

Downgrade to Hold

Down	−4.806***	−1.564***	−3.242***
	(12.56)	(18.77)	(10.26)
Init	−2.718***	−0.579***	−2.139***
	(17.99)	(14.78)	(16.59)
diff	−2.088***	−0.985***	−1.103***
	(7.29)	(14.74)	(5.89)

Downgrade to Sell/Strong Sell

Down	−4.944***	−2.425***	−2.519***
	(3.89)	(8.29)	(2.55)
Init	−4.682***	−1.746***	−2.936***
	(10.03)	(17.56)	(7.97)
diff	−0.262	−0.679***	0.417
	(0.24)	(4.72)	(0.88)

Difference-in-difference results for CARs of revisions made inside vs. outside earnings guidance (EG) windows. The CAR associated with initiations of the corresponding rating is used as the control group. *** and * denote statistical significance at 1% and 10% levels.

from over 0.4 percentage points to more than two percentage points.

Recommendation Surprise and Market Reaction. Loh and Stulz (2009) document that some recommendation revisions are more influential than others. They focus on analyst characteristics and show that some attributes (e.g., experience and reputation) lead to stronger market reactions to revisions. Jegadeesh and Kim (2006) find that stock price responses are stronger following recommendation revisions that are further from the consensus.

Our framework provides a distinct explanation for why some recommendation revisions should have a different market impacts than others. In particular, our model predicts that some recommendations—those where the public-information assessment of the stock valuation is further from a recommendation cutoff—are more surprising than others, as this indicates the analyst's private information must be greater, in order to lead to a recommendation revision. To see this, consider outstanding Hold recommendations for two combinations of observables with different publicly-perceivable valuations. Then a downward revision to Sell conveys more negative information in the scenario with the higher public valuation assessment of the hold, while an upgrade to Buy conveys more positive information for the scenario with the lower public valuation assessment. Consequently, the market CAR responses should be greater following recommendation revisions in these two scenarios. Indeed, to the extent that the findings of Jegadeesh and Kim (2006) are due to larger private information assessments, i.e., greater surprise, our model provides a theoretical rationale for their findings.

We measure the size of recommendation surprise, $\Delta X\beta$, using the difference between the estimate of valuation based on public information $X\beta$ from the *previous month*, and the recommendation friction-adjusted cutoffs corresponding to the revisions/initiations from our full model:

$$\Delta X\beta \equiv \begin{cases} X\beta - (\mu_k + \delta_{k\uparrow}), & \text{upgrades from } k-1 \text{ to } k; \\ X\beta - (\mu_{k+1} - \delta_{k+1,\downarrow}), & \text{downgrades from } k+1 \text{ to } k; \\ X\beta - \mu_k, & \text{new coverage at } k; \end{cases} \tag{8}$$

where k indexes the recommendation level and we omit analyst, firm and time subscripts for simplicity. For initial coverage, we focus on Buy and Strong Buy recommendations because (a) the sample of analysts initiating coverage with a Sell or Strong Sell is too small; and (b) the information content of an initial Hold recommendation is less clear.

$\Delta X\beta$ contains information about the perceived private information content in an analyst's stock recommendation initiation or revision. For instance, given a previous Hold recommendation, the stock valuation V^*, which is the sum of the public assessment $X\beta$ and the private information component u, must breach the friction-adjusted cutoff $(\mu_3 + \delta_{3\uparrow})$ to be upgraded to Buy. A positive $\Delta X\beta$ suggests that an upgrade is widely expected by the market as the public assessment already exceeds the minimum level for triggering such a revision. There should still be a positive market response, reflecting that the market learns that an analyst's private information now exceeds that minimum level. However, the market response should be less than that when $\Delta X\beta$ is very negative, as now an upgrade divulges a positive and potentially very large private information component. Accordingly, the market should react more strongly to such a "surprising" revision. A similar argument holds for downgrades and for new coverage that, for example, is initiated at a buy.

We measure market reactions using three-day market-adjusted CARs. To ensure that the observed return is attributable to a recommendation revision, we exclude recommendations issued within a three-day window around (on and after) quarterly earnings announcement or guidance dates.[18] Our model predicts that the market should react more strongly to those analysts' judgments that imply large private information components. To test this conjecture, we first categorize recommendation changes as upgrades, downgrades and initial coverage. Then, within each category, we sort recommendations into four equal-sized groups based on the sizes of their surprises, as measured by $\Delta X\beta$. Table 12 reports the average size of recommendation surprise, the average 3-day CAR following the recommendation and its standard deviation for each quartile group for the three-tier system.[19] Note that the vast majority of upgrades have $\Delta X\beta < 0$, and almost all downgrades have $\Delta X\beta > 0$, reflecting that our public information measure is calculated based on the *previous* month's public information, and, for example, upgrades typically follow improvement in the public information measures in the current month. The quartile groups are ordered from smallest $\Delta X\beta$ (g1) to largest (g4). The row labeled "$g4 - g1$" shows the difference in average CARs between these two groups.

The least surprising upgrades and initiations are in quartile $g4$, and the least surprising downgrades are in quartile $g1$, where $\Delta X\beta$ is the smallest. The findings in Table 12 strongly indicate that the least surprising recommendations have the smallest market impacts, while the most surprising recommendations have the largest impacts, which are both consistent with the nuanced predictions of our model. The CAR differences between these portfolio quartiles are always substantial and statistically significant. The CAR difference between the most and least surprising quartiles is 1.1 percentage points for upgrades (more surpris-

[18]If we do not discard recommendations issued within an earnings announcement window, CAR impacts become slightly stronger for downgrades and initiations, and slightly weaker for upgrades.

[19]Results for the five-tier system are qualitatively similar, albeit less significant.

ing is good news), −1.33 for downgrades (more surprising is bad news), and 0.52 for Buy initiations (more surprising is good news), and each is statistically significant at the 5% level.

Table 12: Recommendation Surprise vs. Three-day CAR

	Upgrade			Downgrade			Initiation (Buy)		
	Ave $\Delta X\beta$	Ave CAR	Std CAR	Ave $\Delta X\beta$	Ave CAR	Std CAR	Ave $\Delta X\beta$	Ave CAR	Std CAR
$g1$	−1.870	3.751	7.778	1.008	−2.429	8.424	−0.578	2.632	9.606
$g2$	−1.674	3.403	5.542	1.202	−2.861	6.079	−0.398	2.136	4.821
$g3$	−1.478	3.587	6.351	1.379	−2.996	6.929	−0.290	2.175	5.159
$g4$	−0.224	2.638	7.104	2.323	−3.763	7.370	−0.133	2.111	5.681
$g4 - g1$		-1.113^{***}			-1.334^{***}			-0.521^{**}	
		(3.07)			(3.46)			(2.13)	

Equally-weighted quartile portfolios formed by sorting stocks based on the size of recommendation surprise, $\Delta X\beta$. For upgrades and initiations, portfolio $g1$ ($g4$) contains the most (least) surprising ratings, for which the public assessment is below (above) the minimum level for triggering such a revision or initiation. For downgrades, Portfolio $g4$ ($g1$) contains the most (least) surprising ratings. The row "$g4$-$g1$" presents the difference in average CAR between portfolio $g4$ and portfolio $g1$. t-statistics of tests on the equality of means are reported in parentheses. *** denotes statistical significance at the 1% level.

In summary, our framework provides a unique perspective on the market perception of recommendation revisions and coverage initiation. The evidence strongly supports the premise that the market reacts more strongly to decisions by analysts that are more surprising in the context of our model. This suggests both that our model describes how investors believe financial analysts make recommendations, and that investors value the *private* information revealed by analyst's recommendation changes and initiations.

6 Conclusion

We develop a model of how financial analysts formulate recommendations, and show how it captures the rich dynamics in analyst recommendations. Our model incorporates two key features of the recommendation process: (i) analysts acquire information with persistent valuation consequences that the econometrician does not observe, and (ii) analysts revise recommendations reluctantly, introducing frictions to avoid repeatedly revising revisions following small changes in valuation assessments. Our model allows analysts to tailor recommendation revision frictions according to the level of the outstanding recommendation and the direction of a possible revision. Our model nests important existing models as special cases.

Our empirical study reveals that analysts behave quite differently from the "idealized" analyst who has been the focus of existing research. In particular, analysts introduce large recommendation revision frictions to avoid frequent revisions. Our study reveals that publicly-available data on firm and analyst characteristics matters far less for explaining recommendation dynamics than the persistent private information of analysts and their recommendation revision frictions. We find that analysts design the recommendation frictions highly asymmetrically—varying with the recommendation and direction of revision. For example, analysts introduce far smaller frictions "out" of hold recommendations than "into" hold recommendations. Qualitatively, our findings suggest that analysts structure recommendations strategically to generate profitable order flow for their brokerages from their retail clients.

We then document extensive indirect support for our model in (a) durations of recommendation revisions made inside vs. outside earnings announcement and guidance windows, (b) a difference-in-differences analysis of market (CAR) responses to new vs. revised recommendations and inside vs. outside earnings announcement and guidance windows, and (c) CAR responses as a function of the extent to which an analyst's new recommendation or revision is surprising given the extant public information available to the econometrician.

6.1 Appendix: Model Estimation

We use MCMC methods to estimate the model. Following section 3, let $\pi(\theta)$ be the prior distribution of the unknown parameters. If analyst i gave recommendations to firm j for n_{ij} different periods $\{t_{ij0}(s), \cdots, t_{ij*}(s),\ s = 1, \cdots, n_{ij}\}$, then the joint density of θ and r is:

$$\pi(\theta, r) = \pi(\theta) \prod_{j=1}^{J} \prod_{i \in I_j} \prod_{s=1}^{n_{ij}} \left\{ P(R_{ijt_{ij0}(s)} = r_{ijt_{ij0}(s)}) \prod_{t=t_{ij0}(s)+1}^{t_{ij*}(s)} P(R_{ijt} = r_{ijt} | R_{ij,t-1} = r_{ij,t-1}) \right\}, \tag{9}$$

To determine the unconditional and conditional distributions in the above expression, notice that at the beginning of each period, analyst i's initial recommendation $R_{ij,t_{ij0}}$ reflects his valuation $V^*_{ij,t_{ij0}}$. The unobserved residual terms u_{ijt} are modeled by AR processes with parameter ρ,

$$u_{ijt} = \rho u_{ij,t-1} + \varepsilon_{ijt},$$

where ε_{ijt} are i.i.d. $\mathrm{N}(0, \sigma^2)$, and σ^2 is normalized to be 1 for identification purposes. Thus, the variance of u_{ijt} is $\sigma_u^2 = 1/(1 - \rho^2)$, and

$$u_{ij,t_{ij0}} \sim N(0, \sigma_u^2).$$

The probability distribution of the initial recommendation $P(R_{ijt_{ij0}(s)} = r_{ijt_{ij0}(s)})$ can then be calculated from the above normal distribution using the recommendation bins, equation (3). Let $\phi(\cdot)$ be the density function for a $\mathrm{N}(0, 1)$ random variable. Then the density kernel of the initial observation is $\phi\left(u_{ij,t_{ij0}}/\sigma_u\right)$.

After the initial recommendation period, the outstanding recommendation affects analyst's subsequent recommendations at time t and thus the conditional probability $P(R_{ijt} = r_{ijt} | R_{ij,t-1} = r_{ij,t-1})$. Notice that, given $R_{ij,t-1} = r_{ij,t-1}$, we have $R_{ijt} = k$ if

$$\begin{array}{lll} \text{(i) } r_{ij,t-1} > k+1, & \text{and} & \mu_k \le V^*_{ijt} < \mu_{k+1}, \\ \text{(ii) } r_{ij,t-1} = k+1, & \text{and} & \mu_k \le V^*_{ijt} < \mu_{k+1} - \delta_{k+1,\downarrow}, \\ \text{(iii) } r_{ij,t-1} = k, & \text{and} & \mu_k - \delta_{k\downarrow} \le V^*_{ijt} < \mu_{k+1} + \delta_{k+1,\uparrow}, \\ \text{(iv) } r_{ij,t-1} = k-1, & \text{and} & \mu_k + \delta_{k\uparrow} \le V^*_{ijt} < \mu_{k+1}, \\ \text{(v) } r_{ij,t-1} < k-1. & \text{and} & \mu_k \le V^*_{ijt} < \mu_{k+1}. \end{array} \tag{10}$$

Combining (1) and (2) we have:

$$V^*_{ijt} - \rho V^*_{ij,t-1} - \left(X'_{ijt} - \rho X'_{ij,t-1}\right)\beta = \varepsilon_{ijt}.$$

Depending on which group of parameters are studied, it is sometimes convenient to express the above relationship as:

$$\left(V^*_{ijt} - X'_{ijt}\beta\right) - \rho\left(V^*_{ij,t-1} - X'_{ij,t-1}\beta\right) = \varepsilon_{ijt}.$$

Given information up to time t, using equation (10) and noticing that ε_{ijt} are i.i.d. normally distributed, the conditional probability $P(R_{ijt} = r_{ijt} | R_{ij,t-1} = r_{ij,t-1})$ can be obtained from equation (6).

The MCMC estimator using Gibb sampler starts with an initial value $(\theta^{(0)}, V^{(0)})$, and then simulates in turn. Conditional on other parameters and the data, the posterior densities of a subset of parameters can be derived based on the above joint density (9) and given priors. For convenience of conditioning, we divide the vector of parameters θ into 4 groups: (1) β; (2) μ_j, $j = 3, 4, 5$; (3) δ; and (4) ρ. This partition brings a relatively simple form to the conditional posterior densities and makes it more tractable to draw random variables from the conditional distributions. In particular, the conditional distributions of each subset of parameters are given below:

1. The conditional distribution of β is normal. We start with the prior $\beta \sim N(0, I)$. To simplify the simulation, we follow the suggestion of Albert and Chib (1993) and condition on the initial observation. Conditional on the data and other parameters, the conditional distribution of β is given by

$$N(\widehat{\beta}, \widehat{\Sigma}_{\beta}),$$

where

$$\begin{aligned} \widehat{\beta} \;\; = \;\; & \left[\sum_{j=1}^{J} \sum_{i \in I_j} \prod_{s=1}^{n_{ij}} \prod_{t=t_{ij0}(s)+1}^{t_{ij*}(s)} \left(X_{ijt} - \rho X_{ij,t-1}\right) \left(X'_{ijt} - \rho X'_{ij,t-1}\right) \right]^{-1} \\ & \sum_{j=1}^{J} \sum_{i \in I_j} \prod_{s=1}^{n_{ij}} \prod_{t=t_{ij0}(s)+1}^{t_{ij*}(s)} \left(X_{ijt} - \rho X_{ij,t-1}\right) \left(V^*_{ijt} - \rho V^*_{ij,t-1}\right), \end{aligned}$$

and variance (inverse precision)

$$\Sigma_{\widehat{\beta}} = \left[\sum_{j=1}^{J} \sum_{i \in I_j} \prod_{s=1}^{n_{ij}} \prod_{t=t_{ij0}(s)+1}^{t_{ij*}(s)} \left(X_{ijt} - \rho X_{ij,t-1}\right) \left(X'_{ijt} - \rho X'_{ij,t-1}\right) \right]^{-1},$$

where $X_{ij0} = 0$.

2. The conditional distribution of ρ is a truncated normal. We start with the prior $N(0.5, 1) I(|\rho| < 1)$. Conditional on the data and other parameters, ρ is normally distributed

with mean $\widehat{\rho}$, and variance $\widehat{\sigma}_\rho^2$, truncated by $|\rho| < 1$, i.e., $\rho \sim N(\widehat{\rho}, \Sigma_\rho) \cdot I(|\rho| < 1)$, where

$$\begin{aligned} \widehat{\rho} &= \left[\sum_{j=1}^{J} \sum_{i \in I_j} \prod_{s=1}^{n_{ij}} \prod_{t=t_{ij0}(s)+1}^{t_{ij*}(s)} \left(V_{ij,t-1}^* - X_{ij,t-1}' \beta \right)^2 \right]^{-1} \\ & \quad \sum_{j=1}^{J} \sum_{i \in I_j} \prod_{s=1}^{n_{ij}} \prod_{t=t_{ij0}(s)+1}^{t_{ij*}(s)} \left(V_{ij,t-1}^* - X_{ij,t-1}' \beta \right) \left(V_{ij,t}^* - X_{ij,t}' \beta \right), \end{aligned}$$

and

$$\widehat{\sigma}_\rho^2 = \left[\sum_{j=1}^{J} \sum_{i \in I_j} \prod_{s=1}^{n_{ij}} \prod_{t=t_{ij0}(s)+1}^{t_{ij*}(s)} \left(V_{ij,t-1}^* - X_{ij,t-1}' \beta \right)^2 \right]^{-1}.$$

3. The conditional density of $\delta_{k\uparrow}$ is a uniform distribution. Given the data, other parameters ρ, β, μ, and other elements in δ, the conditional density of $\delta_{k\uparrow}$ is proportional to $\pi(\theta, r)$. The bounds of the uniform distribution can be derived based on the following information: (1) at the starting period, the stickiness parameter does not enter, and (2) at other periods, the above likelihood is non-zero when: (i) $\mu_k - \delta_{k\downarrow} \le V_{ijt}^* < \mu_{k+1} + \delta_{k+1,\uparrow}$ if $R_{ijt} = k$ and $R_{ij,t-1} = k$; (ii) $\mu_k \le V_{ijt}^* < \mu_{k+1} - \delta_{k+1,\downarrow}$ if $R_{ijt} = k$ and $R_{ij,t-1} = k+1$; and (iii) $\mu_k + \delta_{k\uparrow} \le V_{ijt}^* < \mu_{k+1}$ if $R_{ijt} = k$ and $R_{ij,t-1} = k-1$, (other cases do not depend on the friction parameters.) Using the above information, we obtain that $\delta_{k\uparrow}$ is uniformly distributed on

$$[\delta_{k\uparrow low}, \delta_{k\uparrow up}],$$

where

$$\begin{aligned} \delta_{k\uparrow low} &= \max_{i,j, \text{ and } t \neq t_{ij0}(s)} \left\{ V_{ijt}^* - \mu_k : R_{ijt} = R_{ij,t-1} = k-1 \right\}, \\ \delta_{k\uparrow up} &= \min_{i,j, \text{ and } t \neq t_{ij0}(s)} \left\{ \left(V_{ijt}^* - \mu_k \right) \wedge \left(\mu_k - \mu_{k-1} \right) : R_{ijt} = k,\ R_{ij,t-1} = k-1 \right\}. \end{aligned}$$

Similarly, given the data and other parameters, $\delta_{k\downarrow}$ is uniformly distributed on

$$[\delta_{k\downarrow low}, \delta_{k\downarrow up}],$$

where

$$\begin{aligned} \delta_{k\downarrow low} &= \max_{i,j, \text{ and } t \neq t_{ij0}(s)} \left\{ \mu_k - V_{ijt}^* : R_{ijt} = k = R_{ij,t-1} \right\}, \\ \delta_{k\downarrow up} &= \min_{i,j, \text{ and } t \neq t_{ij0}(s)} \left\{ \left(\mu_k - V_{ijt}^* \right) \wedge \left(\mu_k - \mu_{k-1} \right) : R_{ijt} = k-1,\ R_{ij,t-1} = k \right\}. \end{aligned}$$

4. The conditional density of μ_k given the data and other parameters ρ, β, δ, and $\mu_l \neq \mu_k$, is a uniform distribution on the interval $[\mu_{k,low}, \mu_{k,up}]$. The lower bound and upper bound

can be derived in a similar way as δ. In particular, the lower bound is

$$\mu_{k,low} = \max\left\{\mu_{k-1}, \max_{s}\left[\ V^*_{ijt_{ij0}(s)}\ |R_{ijt_{ij0}(s)} = k-1\right], \mu_{k,l1}, \mu_{k,l2}, \mu_{k,l3}, \mu_{k,l4}\right\},$$

where

$$\begin{aligned}
\mu_{k,l1} &= \max_{t\neq t_{ij0}(s)}\left[\ V^*_{ijt} - \delta_{k+1,\uparrow}|\ R_{ijt} = R_{ij,t-1} = k\right], \\
\mu_{k,l2} &= \max_{t\neq t_{ij0}(s)}\left[\ V^*_{ijt} + \delta_{k+1\downarrow}\ |R_{ijt} = k, R_{ij,t-1} = k+1\right], \\
\mu_{k,l3} &= \max_{t\neq t_{ij0}(s)}\left[\ V^*_{ijt}\ |R_{ijt} = k-1, R_{ij,t-1} > k\right], \\
\mu_{k,l4} &= \max_{t\neq t_{ij0}(s)}\left[\ V^*_{ijt}\ |R_{ijt} = k-1, R_{ij,t-1} \leq k-2\right],
\end{aligned}$$

and, similarly, for the upper bound:

$$\mu_{k,up} = \min\left\{\mu_{k+1}, \min_{s}\left[\ V^*_{ijt_{ij0}(s)}\ |R_{ijt_{ij0}(s)} = k\right], \mu_{k,u1}, \mu_{k,u2}, \mu_{k,u3}, \mu_{k,u4}\right\}.$$

where

$$\begin{aligned}
\mu_{k,u1} &= \min_{t\neq t_{ij0}(s)}\left[V^*_{ijt} + \delta_{k\downarrow}\ |R_{ijt} = R_{ij,t-1} = k\right], \\
\mu_{k,u2} &= \min_{t\neq t_{ij0}(s)}\left[\ V^*_{ijt}\ -\delta_{k\uparrow}|R_{ijt} = k, R_{ij,t-1} = k-1\right], \\
\mu_{k,u3} &= \max_{t\neq t_{ij0}(s)}\left[\ V^*_{ijt}\ |R_{ijt} = k, R_{ij,t-1} \geq k+1\right], \\
\mu_{k,u4} &= \max_{t\neq t_{ij0}(s)}\left[\ V^*_{ijt}\ |R_{ijt} = k, R_{ij,t-1} < k-1\ ,\right].
\end{aligned}$$

5. The conditional distributions of the latent variables V^*_{ijt} are truncated normal. Giving our valuation model in Section 2, conditional on data and other information, the conditional distributions of V^*_{ijt} at the initial period $t_{ij0}(s)$ are given by

$$V^*_{ijt_{ij0}(s)} \sim N\left(\underline{X}'_{ij}\underline{\beta}, \frac{1}{1-\rho^2}\right), \text{ truncated by } [\mu_{R_{ijt_{ij0}(s)}}, \mu_{R_{ijt_{ij0}(s)}+1}].$$

For subsequent periods, i.e., $t \in \{t_{ij0}(s)+1, \cdots, t_{ij*}(s), s = 1, \cdots, n_{ij}\}$, notice that $V^*_{ijt} = \rho V^*_{ij,t-1} + \left(X'_{ijt} - \rho X'_{ij,t-1}\right)\beta + \varepsilon_{ijt}$, the conditional distributions of V^*_{ijt} are truncated normals with means $\rho V^*_{ij,t-1} + \left(X'_{ijt} - \rho X'_{ij,t-1}\right)\beta$ and unit variance, truncated at $[\mu_{t,low}, \mu_{t,upp}]$, where

$$\begin{aligned}
\mu_{t,low} &= 1(R_{ijt} = R_{ij,t-1}+1)(\mu_{R_{ijt}} + \delta_{\uparrow}) + 1(R_{ijt} = R_{ij,t-1})(\mu_{R_{ijt}} - \delta_{\downarrow}) \\
&\quad +1(R_{ijt} \leq R_{ij,t-1} - 1 \text{ or } R_{ijt} > R_{ij,t-1}+1)\mu_{R_{ijt}}\ ,
\end{aligned}$$

$$\mu_{t,upp} = 1(R_{ijt} = R_{ij,t-1})(\mu_{R_{ijt}+1} + \delta_{\uparrow}) + 1(R_{ijt} = R_{ij,t-1} - 1)(\mu_{R_{ijt}+1} - \delta_{\downarrow})$$
$$+1(R_{ijt} < R_{ij,t-1} - 1 \text{ or } R_{ijt} \geq R_{ij,t-1} + 1)\mu_{R_{ijt}+1} \, .$$

Fix a draw q, and denote the conditional distribution of, say, V conditional on θ as $p(V|\theta)$, where the conditioning on X and R is suppressed.

1. Draw $V^{(q)}$ from $p(V|\theta^{(q-1)})$, i.e., $p(V|\beta^{(q-1)}; \mu_j^{(q-1)}; \delta^{(q-1)}; \rho^{(q-1)})$
2. Draw $\beta^{(q)}$ from $p(\beta|V^{(q)}; \mu_j^{(q-1)}; \delta^{(q-1)}; \rho^{(q-1)})$
3. Draw $\mu_j^{(q)}$ $j = 3, 4, 5$, from $p(\mu|V^{(q)}; \beta^{(q)}; \delta^{(q-1)}; \rho^{(q-1)})$. This is done by the following 3 steps:
   (a) draw $\mu_3^{(q)}$ from $p(\mu_3|V^{(q)}; \beta^{(q)}; \mu_4^{(q-1)}, \mu_5^{(q-1)}; \delta^{(q-1)}; \rho^{(q-1)})$.
   (b) draw $\mu_4^{(q)}$ from $p(\mu_4|V^{(q)}; \beta^{(q)}; \mu_3^{(q)}, \mu_5^{(q-1)}; \delta^{(q-1)}; \rho^{(q-1)})$.
   (c) draw $\mu_5^{(q)}$ from $p(\mu_5|V^{(q)}; \beta^{(q)}; \mu_3^{(q)}, \mu_4^{(q)}; \delta^{(q-1)}; \rho^{(q-1)})$.
4. Draw $\delta^{(q)}$ from $p(\delta|V^{(q)}; \beta^{(q)}; \mu_j^{(q)};\ \rho^{(q-1)})$. For example, if $\delta = (\delta_{\uparrow}, \delta_{\downarrow})$, this is done by the following 2 steps:
   (a) Draw $\delta_{\uparrow}^{(q)}$ from $p(\delta_{\uparrow}|V^{(q)}; \beta^{(q)}; \mu_j^{(q)}; \delta_{\downarrow}^{(q-1)}; \rho^{(q-1)})$.
   (b) Draw $\delta_{\downarrow}^{(q)}$ from $p(\beta|V^{(q)}; \beta^{(q)}; \mu_j^{(q)}; \delta_{\uparrow}^{(q)}; \rho^{(q-1)})$.
5. (5) Draw $\rho^{(q)}$ from $p(\rho|V^{(q)}; \beta^{(q)}; \mu_j^{(q)}; \delta^{(q)})$
6. Set $q = q + 1$, and repeat the above steps.

In practice, we discard the first M draws (in our empirical analysis, we set $M = 50,000$), and the simulated values of $(\theta^{(q)}, V^{(q)})$ from $q = M + 1, \cdots, M + Q$, can be regarded as an approximate simulated sample. The posterior expectation of a function of the parameters, $h(\theta)$, can then be estimated by the sample average

$$\frac{1}{Q} \sum_{q=M+1}^{M+Q} h(\theta^{(q)}).$$

7 References

Agrawal, A. and M. Chen, 2008, "Do Analyst Conflicts Matter? Evidence from Stock Recommendations." Journal of Law and Economics, 51(3), 503-537.

Bagnoli, Mark, M. Clement, M. Crawley, and Susan Watts, 2009, The Profitability of Analysts' Stock Recommendations: What Role Does Investor Sentiment Play?, SSRN working paper, http://ssrn.com/.

Barber, Brad, Reuven Lehavy, Maureen F. McNichols and Brett Trueman, 2001, "Can Investors Profit From The Prophets? Security Analyst Recommendations and Stock Returns", *Journal of Finance* 56, 531-563.

Barber, Brad, Reuven Lehavy, Maureen F. McNichols and Brett Trueman, 2005, "Buys, Holds, and Sells: The Distribution of Investment Banks' Stock Ratings and the Implications for the Profitability of Analysts' Recommendations" SSRN working paper.

Bernhardt, D., M. Campello and E. Kutsoati, 2006, "Who Herds?", with E. Kutsoati and M. Campello, *Journal of Financial Economics*, 657-675.

Boni, Leslie and Kent L. Womack, 2006, "Analysts, industries, and price momentum", *Journal of Financial & Quantitative Analysis* 41, 85-109.

Brandt, Michael W., Runeet Kishore, Pedro Santa-Clara, and Mohan Venkatachalam, 2008, "Earnings Announcements are Full of Surprises", working paper, Duke University.

Cowles, M. 1996). "Accelerating Monte Carlo Markov chain convergence for cumulative-link generalized linear models", *Statistics and Computing* 6, 101-111.

Conrad, Jennifer, Bradford Cornell, Wayne R. Landsman and Brian R. Rountree, 2006, "How Do Analyst Recommendations Respond to Major News?", *Journal of Financial and Quantitative Analysis*, 41:25-49 Cambridge University Press.

Diether, K., C. Malloy, and A. Scherbina, 2002, "Differences of opinion and the cross section of stock returns", *Journal of Finance*, 57, 2113-2141.

Evett, I.W., 1991, "Implementing Bayesian Methods in Forensic Science", paper presented at the Fourth Valencia International Meeting on Bayesian Statistics.

Fama, E., French, K., 1993, "Common risk factors in the returns on stocks and bonds," *Journal of Financial Economics* 33, 3-56.

Geweke, J., 1992, Evaluating the accuracy of sampling-based approaches to calculating posterior moments. In Bayesian Statistics 4 (ed. JM Bernado, JO Berger, AP Dawid and AFM Smith). Clarendon Press, Oxford, UK.

Hackethal A., R. Inderst, and S. Meyer, 2010, "Trading on Advice," working paper, University of Frankfurt.

Ivkovic, Zoran, Jegadeesh, Narasimhan, 2004, "The timing and value of forecast and recommendation revisions", *Journal of Financial Economics* 73, 433-463.

Jegadeesh, Narasimhan, Joonghyuk Kim, Susan Krische, and Charles Lee, 2004, "Analyzing the Analysts: When do recommendations add value?", *Journal of Finance* 59, 2004, 1083-1124.

Jegadeesh, Narasimhan, and Woojin Kim, 2010 "Do analysts herd? An analysis of recommendations and market reactions." *Review of Financial Studies* 23, 901-937.

Kadan, O., L. Madureira, R. Wang, and T. Zach, 2009, "Conflicts of Interest and Stock Recommendations — the Effect of the Global Settlement and Related Regulations," *Review of Financial Studies*, 22, 4189-4217.

Kass, Robert E. and Raftery, Adrian E., 1995, "Bayes Factors", *Journal of American Statistical Association*, 90, 773-795.

Kecskés, Ambrus, Roni Michaely and Kent L. Womack, 2010, "What Drives the Value of Analysts' Recommendations: Earnings Estimates or Discount Rate Estimates?", Tuck School of Business Working Paper No. 2009-67.

Lakonishok, Joseph, Andrei Shleifer and Robert W. Vishny, 1994, "Contrarian Investment, Extrapolation, and Risk", *Journal of Finance* 49, 1541-1578.

Ljungqvist, A., F. Marston and W.J. Wilhelm, 2006, "Competing for Securities Underwriting Mandates: Banking Relationships and Analyst Recommendations," *Journal of Finance*

Ljungqvist, Alexander, Felicia Marston, Laura T. Starks, Kelsey D. Wei and Hong Yan, 2007, "Conflicts of interest in sell-side research and the moderating role of institutional investors", *Journal of Financial Economics* 85, 420-456.

Loh, R., and R. M. Stulz, 2011, "When are Analyst Recommendation Changes Influential?", *Review of Financial Studies* 24, 593-627.

Malmendier, Ulrike and Devin Shanthikumar, 2007, "Are small investors naive about incentives?", *Journal of Financial Economics* 85 457-489.

Muller, Gernot and Claudia Czado, 2005, "An autoregressive Ordered Probit Model with

Application to High Frequency Financial Data". *Journal of Computational & Graphical Statistics* 14, 320-338.

Raedy, J., P. Shane and Y. Yang, 2006, Horizon-dependent underreaction in financial analysts' earnings forecasts". *Contemporary Accounding Research* 23, 291-322.

Stickel, S. 1995, "The Anatomy of Performance of Buy and Sell Recommendations", *Financial Analysts Journal* 51, 25-39.

Womack, K.L., 1996, "Do Brokerage analysts' recommendations have investment value?" *Journal of Finance* 51, 137-167.

Zhang, X., 2006, "Information uncertainty and analyst forecast behavior", *Contemporary Accounting Research* 23, 565-590.