How often do managers withhold information?*

Jeremy Bertomeu, Paul Ma, and Iván Marinovic

Abstract

This study structurally estimates the model of voluntary disclosure with uncertain information endowment of Dye (1985) in the context of quarterly earnings forecasts of US large cap firms. We find that managers strategically withhold information on average one out of seven quarters. Conditional on not observing a forecast, the probability of strategic withholding is 14%, in which case the concealed information falls in the bottom quartile of the distribution. The private information of managers at a disclosure date explains 27% of the variation in realized earnings. Disclosure frictions tend to be sticky, and their frequency varies across industries. Firms with less precise information are more likely to selectively withhold information. Comparing the estimated volatility of managerial beliefs to the volatility of realized earnings, we show that expectations exhibit excess volatility for about two thirds of the sample firms, which we interpret as a sign of managerial overconfidence. Our study is the first to provide large sample estimates of the limits to unravelling.

Keywords: voluntary disclosure, disclosure frictions, structural estimation
JEL Classification: D72, D82, D83, G20

*Jeremy Bertomeu is an assistant professor at Baruch College, City University of New York, 1 Bernard Baruch Way, New York, NY 10010. Email: jeremy.bertomeu@baruch.cuny.com. Paul Ma is an assistant professor at the Carlson School of Management, University of Minnesota, 321 19th Avenue South, Minneapolis, MN 55455. Email: paulma@umn.edu. Iván Marinovic is an assistant professor at Stanford Graduate School of Business, Stanford University, 655 Knight Way, Stanford, CA 94305. Email: imvial@stanford.edu. We are grateful for comments from Santiago Bazdresch, Alan Benson, Anne Beyer, William Greene, David Larcker, Chandra Kanodia, Colleen Manchester, Alejandro Molnar, Raj Singh, Aaron Sojourner, John Shoven, Stephen Terry, seminar participants at Duke and the University of Chicago. All errors are our own.

Persuasion theory suggests that, if frictions prevent costless communication of private information, an informed sender will strategically withhold information (Dye 1985; Fishman and Hagerty 1990; Shavell 1994; Shin 1994, 2003). The law literature offers many instances in which material information was voluntarily omitted, with cases of successful litigation in the pharmaceutical, food, tobacco and automobile industries (see Rabin 2001; Mello, Rimm and Studdert 2003; McClellan 2008; Govindaraj, Lee and Tinkelman 2007). These cases appear to be consistent with evidence, in experimental settings, that informed senders conceal unfavorable information as a result of disclosure frictions (Dickhaut, Ledyard, Mukherji and Sapra 2003; Jin, Luca and Martin 2015).

However, to our knowledge, most of the field evidence about strategic withholding focuses on unusual events, as in the large body of evidence regarding safety issues. Such events are usually easier to observe ex-post, from the documentation collected during lawsuits. Lawsuits are somewhat problematic for a field test of voluntary theory because they imply that disclosures are not strictly voluntary. In this study, we structurally estimate the voluntary disclosure theory of Dye (1985) and Jung and Kwon (1988), hereafter DJK, using quarterly earnings forecasts. We rely primarily on this theory as our baseline, because it is one of the most widespread theory and continues to inspire both theoretical and empirical research (Guttman, Kremer and Skrzypacz 2014; Ben-Porath, Dekel and Lipman 2014).

Earnings forecasts are an ideal setting to estimate voluntary disclosure theory. In practice, firms occasionally make voluntary disclosure of next-quarter earnings. These forecasts are protected against most lawsuits by safe harbor provisions (see *Safe Harbor for Forward Looking Statements, Concept Release and Notice of Hearing, October 13, 1994*).[1] Earnings forecasts offer rich panel data over homogenous quantitative forecasts that spans across industries, and in which we can match each forecast to the realized

[1]Some disclosures are subject to a duty to update (or duty to correct), under section 10(b) of the Securities and Exchange Act of 1934, but this requirement is interpreted as applying to material facts that may have changed - not to quantitative quarterly earnings forecast.

forecasted earnings in the next quarter and to a market expectation at the time forecast is made, using the pre-forecast *First Call* analyst consensus.

Earnings forecasts are also of great stand-alone interest in capital market research. Prior theoretical research conjectures that voluntary disclosure is a primary channel through which the classic lemons problem is resolved. Empirically, Ball and Brown (1968) find that most of the information becomes known to the market before the release of mandated financial statements. Considering earnings forecasts, Beyer, Cohen, Lys and Walther (2010) document that these forecasts account for as much as 16% of quarterly stock return variance, more than earnings announcements, filings with the Securities and Exchange Commission and analyst forecasts combined.

In this paper, our objective is to quantify the limits to voluntary disclosure, by structurally estimating disclosure frictions and analyzing its implications for strategic withholding and the range of concealed information. In the DJK model, a manager may be privately informed about the value of a traded firm and, if so, decides whether to disclose or withhold the information, aiming to maximize the *current* market value of the firm. Absent any friction, disclosure strategies unravel to full disclosure because a manager with the best prospects among other withholding managers will always be better-off disclosing (Grossman and Hart, 1980; Milgrom, 1981; Grossman, 1981). To break unraveling, the model conjectures that there is a random probability of a friction, unobservable to outsiders, which makes the manager either unable or unwilling to disclose *regardless of the observed private information.*[2] As long as the probability of the friction is non-zero, even a manager that is not subject to the friction will selectively withhold some unfavorable information.

The disclosure friction in DJK is an abstract mechanism that has many possible

[2]In DJK, the working assumption is that the probability of the friction is not a function of the true information - it is generally used in the literature as a synonym for this model and is analogue to the fixed disclosure assumption in costly disclosure models (Jovanovic 1982). Another key assumption is that insiders cannot trade strategically; as is illustrated in the market microstructure literature, any information that affects price discovery will greatly complicate trading strategies (Huddart, Hughes and Levine 2001) and, in our model, would feed back into complex disclosure strategies.

interpretations. In the conventional interpretation proposed by Dye (1985), managers may be uninformed, a fact that they cannot credibly disclose. Alternatively, managers might occasionally face prohibitive proprietary costs if they disclose in advance. For our application, which relies on relatively short-term horizons (a forecast about next quarter earnings), one may consider a third interpretation – uncertainty about managerial objectives – in the spirit of Fischer and Verrecchia (2000). When the friction is *not* present, the manager has a short-term focus and selectively discloses to maximize the current stock price; by contrast, when the friction is present, the manager has a pure long-term focus and, assuming some small cost of disclosure, prefers to withhold his forecast.

The main prediction of DJK is that the distribution of forecasts is a truncated version of the distribution of managers' beliefs, because managers withhold a fraction of their bad news. The truncation is reminiscent of that arising in the classic selection models popularized by Tobin (1958) and Heckman (1979) in the context of labor markets, where the distribution of wage offers of employed workers is a truncation of the unconditional distribution of wage offers. In a standard selection models, a common approach is to assume that the truncation point is zero (see Amemiya 1984, p. 7). In our setting, it is crucial to estimate the truncation point (that is, the disclosure threshold) which we then use, in conjunction with the structural equations of the disclosure game, to infer the structural process of the disclosure friction.

Another challenge in our empirical implementation is to accommodate stickiness in disclosure policies (see Graham, Harvey and Rajgopal 2005). To make the theory amenable to empirical testing, we augment DJK with time-series autocorrelation in the friction. This means that the friction follows a hidden Markov process. If the friction is sticky, a disclosure decision today affects the evolution of market beliefs about the friction, thereby affecting the evolution of disclosure incentives. Infrequent disclosures in the past may lead investors to believe that the manager is likely to experience frictions in the future and react more favorably to the absence of disclosure.

We estimate the model at the firm level using a maximum likelihood approach. We recover the following parameters: the sequence of disclosure threshold; the probability and persistence of the disclosure friction, and the quality of the manager's private information. Based on these estimates, we then compute the average probability of strategic withholding.

We first document that managers possess significant private information: roughly, managers' private information accounts for a quarter of the variation in earnings innovations. On average, managers strategically withhold information about one out of then quarters. Frictions occur, on average, only once every two quarters but market pressure, in the form of negative expectations conditional on withholding induces managers to strategically withhold only significantly unfavorable news in order to avoid the market penalty. Taken together, these results imply that managers conceal the bottom quartile of the distribution.

We also test for managerial biases. According to elementary statistics, the variance of a Bayesian expectation of a random variable cannot be greater than the variance of that random variable.[3] This prediction cannot be tested directly because the variance of manager expectations is not directly observable - we only observe forecasts when the expectation is favorable. Using the model's equilibrium characterization, however, we are able to recover the location of the truncation which in turn allows us to compute the variance of manager expectations. If managers make their predictions rationally, this variance will be lower than the variance of the underlying earnings. By contrast, if managers are over-confident in their own private information, the variance of the prediction could be higher than that of earnings.

We construct a test based on this argument and reject Bayesian rationality for about

[3]This bound is also commonly used in tests of excess price volatility (Shiller 1980; LeRoy and Porter 1981). It is recently used in the context of analysts' forecasts by Lundholm and Rogo (2014); these are easier to examine because they are less likely to exhibit a truncation. They show that, like for management forecasts, there is excess volatility of forecasts for about half of the firms.

60% of the sample firms. Indeed, for about two thirds of the sample firms, manager forecasts exhibit excess volatility relative to the volatility of earnings. For the median firm, the variance of manager expectations is 25% higher than if managers were consistent with Bayes' rule. Taken together, these estimates provide evidence of a behaviorial bias in managers' disclosures.

Related literature. Our paper offers a structural estimation of a formal model of communication using field data. To our knowledge, few current studies provide structural estimates of models of strategic communication. Beyer, Guttman and Marinovic (2014) estimate a dynamic misreporting model. In their model, the manager can bias a mandatory report for a cost (Goldman and Slezak, 2006; Kedia and Philippon, 2009; Acharya and Lambrecht, 2011) but the market cannot recover the true signal because of noise in the accounting system. Zakolyukina (2014) and Terry (2014) estimate structural models in which agents consider the dynamic consequences of manipulation. Zakolyukina (2014) estimates her model using observed accounting violations, recovering current choice of manipulation as a trade-off between current benefits and an increase in the long-term probability of an accounting restatement. Terry (2014) focuses on the effect of misreporting on firm's dynamic investment policy, measuring that reporting incentives appear to cause significant distortions to investment. To our knowledge, the only study that implements a test of a voluntary forecasting model is Chen and Jiang (2006); they focus on how analysts weight their information when making forecasts.[4]

There is also an extensive theoretical literature on voluntary disclosure drawing on the original model of Dye (1985). One of the predictions of this theory is the use of sanitization strategies, in which unfavorable information is withheld Shin (1994, 2003). The DJK model, which we estimate here, is one form of sanitization but the literature has shown that this can also be affected by various other factors. For example, in Hagerty and Fish-

[4]For other examples of structural estimation using field data, see Erdem and Keane (1996), Bajari and Hortacsu (2003) or Gentzkow and Shapiro (2010).

man (1989) and Bushman (1991), the amount of information released is a function of the price discovery in the market. Fishman and Hagerty (1990) show that some mandatory disclosure will affect voluntary disclosure, in a model where the manager might cherry-pick across multiple dimensions of information. Dye and Sridhar (1995) and Acharya, Demarzo and Kremer (2011) show that, given correlation between firms' information, firms' disclosure can be positively correlated to other firms' disclosures. While we have left these factors aside here, our methodology may, in the future, allow researchers to condition the estimation on other observable feature.

1 Theoretical development

1.1 Baseline Model

We briefly summarize the model of Dye (1985) and Jung and Kwon (1988), and then extend it to a multi-period setting. The dynamic model will be the basis for most of the empirical analysis developed in Section 2.

There is a single firm and a mass of risk neutral investors (the market). The firm's manager is privately informed about the firm's prospects. Investors believe that there is a probability p that the manager is subject to a friction under which he is unable to make a disclosure – e.g. because he is uninformed or lacks incentives to disclose but for reasons that are independent of his private information.

With probability $1 - p$, the manager is not subject to the friction, and may either truthfully disclose a forecast x about the firm's future prospects v, or withhold that information altogether. In the standard DJK model, the manager does not manipulate the information but, in our estimation procedure, we will allow for possible disclosure biases.[5] When not subject to the friction, the manager makes his disclosure choice d,

[5]Dye (1985) and Shin (2003) assume that there exists an ex-post verification technology. In our setting, realized earnings tend to be very close to the forecasts: forecast errors are on average about 10% of earnings per share. Stocken (2000) suggests that, even if this monitoring is imperfect, the repeated

seeking to maximize the current period's market price. The firm's stock price is an increasing function of the manager's information, x.

The equilibrium has a simple structure. In the absence of the friction, the manager discloses his information if and only if $x \geq y$. Investors update their beliefs in accordance with Bayes' rule and the manager's disclosure strategy. For the manager to be indifferent between disclosing and not disclosing his information, when $x = y$, the following condition must hold:

$$y = \frac{pE[x]}{p + (1-p)\Pr(x \leq y)} + \frac{(1-p)\Pr(x \leq y)E[x|x \leq y]}{p + (1-p)\Pr(x \leq y)}, \tag{1}$$

where the left hand side is the market belief induced by a disclosure at the margin, and the right hand side is the market belief when there is no disclosure. Jung and Kwon (1988) prove that there is a unique threshold y that solves equation (1). They further show that the threshold increases in the probability of the friction p, and decreases when the distribution of v experiences a mean preserving increasing spread.

We now extend this model to a multi-period setting, using index $t = 1, 2, ..., T$ for each period. We omit the firm index i, since we carry out the estimation at the firm level.

There are three dates (see Figure 1). First, investors observe p_t, which represents the perceived probability that the manager is subject to a disclosure friction in period t. For the moment, we defer the discussion of how p_t is updated and instead treat p_t as exogenous. In addition to p_t investors observe μ_t and update their perceptions of the manager's information x_t and the firm's earnings v_{t+1} as given by

$$(x_t, v_{t+1})\Big|_{\mu_t} \sim N\left(\begin{bmatrix}\mu_t \\ \mu_t\end{bmatrix}, \Sigma \equiv \begin{bmatrix}\sigma^2 & \rho\sigma\varphi \\ \rho\sigma\varphi & \varphi^2\end{bmatrix}\right), \tag{2}$$

with p.d.f. f_t and c.d.f. F_t. In this model, σ represents the dispersion of the managers' expectations and can be interpreted as the quality of their information: a manager with

nature of the game can nevertheless enforce truthful reporting. In the estimation, we still allow the reports to be biased (a parameter we estimate) to correct for a fully-separating biased equilibrium under costly misreporting (Dye 1988; Stein 1989; Kartik, Ottaviani and Squintani 2007).

more precise information will have more disperse beliefs (Ganuza and Penalva, 2010). This parameter may capture both the quality of the firm's accounting system, or some aspects of the firm's underlying business model. For example, in some industries, the manager's uncertainty about the next period earnings may be resolved early during the quarter, whereas in some other industries the uncertainty may be resolved late.

The parameter φ, by contrast, represents the dispersion of earnings, capturing the market's uncertainty about the firm's future performance. More volatile industries are characterized by a larger φ. Finally, ρ represents the correlation between x_t and v_{t+1} also reflecting the quality of the manager's information. In the baseline setting, we assume that the manager is rational so his beliefs are given by $x_t = E_t[v_{t+1}]$. This allows us to write $v_{t+1} = x_t + \varepsilon_t$, where ε_t is white noise. Consequently, the variance-covariance matrix reduces, in this case, to $\Sigma = \begin{bmatrix} \sigma^2 & \sigma^2 \\ \sigma^2 & \varphi^2 \end{bmatrix}$, that is, $\rho = corr_t(x_t, v_{t+1}) = \sigma/\varphi$.

Second, if not subject to the friction, the manager chooses whether or not to disclose his information. We denote by $d_t \in \{\oslash, x_t\}$ the manager's disclosure decision in period t. As mentioned above, the manager makes this decision aiming to maximize the firm's current stock price which strictly increases in x_t. This captures, for example, a manager who is willing to accelerate the effect of news before the next-quarter earnings announcement.[6]

Third, the earnings v_{t+1} are realized and publicly observed. The markets uses this information to update the probability that the manager concealed information in the previous period, and the probability that the manager experiences the friction in the next period.

[6]The DJK model is, in essence, a model of managerial short-termism. As in Shin (2003), while our model has multiple periods and features learning over time, we assume that the manager maximizes current stock price. Short-termism is at the core of disclosure theory since, without any short-term motive, the manager could wait for realized cash flows. One may criticize this approach on the grounds that the reality is probably more complex, with managers maximizing prices over time. This noted, we believe that the DJK model provides a useful parsimonious benchmark. Unfortunately, forward-looking motives would require us to place much more structure on intertemporal preferences and, because of the rational pricing function, imply a computationally difficult problem. We are not aware of any theoretical treatment of a fully dynamic (infinite horizon) disclosure model with forward-looking motives.

Figure 1. Time Line

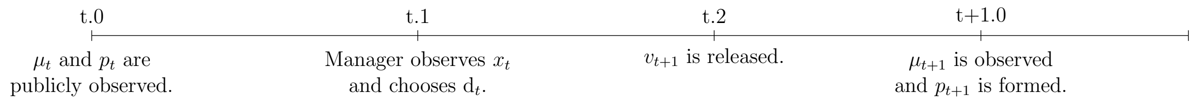

Preliminaries In the absence of the friction, the manager is myopic. Therefore, for a given p_t his strategy is characterized in equation (7) of Jung and Kwon (1988). The manager will issue a forecast if and only if $x_t \geq y_t$ where y_t is given by

$$p_t(\mathbb{E}_t(x_t) - y_t) = (1 - p_t) \int_{-\infty}^{y_t} F_t(x)dx. \tag{3}$$

Using standard properties of Normal distributions, Equation (3) reduces to

$$-\frac{p_t}{1-p_t}(y_t - \mu_t) = \int_{-\infty}^{y_t} \Phi(\frac{x-\mu_t}{\sigma})dx,$$

where Φ is the c.d.f. of the standard Normal. With a change of variables, this equation simplifies to

$$-\frac{p_t}{1-p_t}z_t = \int_{-\infty}^{z_t} \Phi(x)dx \tag{4}$$

where $z_t \equiv \frac{y_t - \mu_t}{\sigma}$. Equation (4) defines an implicit function $z_t = Z(p_t)$ that maps the probability of a friction p_t to a standardized disclosure threshold z_t. The probability of observing a forecast in period q_t is given by

$$q_t = (1 - p_t)\,\Phi(-Z(p_t)). \tag{5}$$

So, two events must happen at the same time to trigger a disclosure: the absence of a friction, and the realization of favorable information. Equation (5) forms the basis of our empirical strategy, as it describes the relation between the (observable) frequency of forecast, q_t, and the probability of the friction, p_t. That is, it implicitly defines a function $P(q_t)$ such that, for any q_t, where $p_t = P(q_t)$ would be a solution to the above equation.

Conveniently, under normality, the function $P(.)$ does not require knowledge of μ_t or σ_t, and the next claim guarantees that p_t is identified.

Claim 1 *The function $P(.)$ is a strictly decreasing, continuous, invertible bijection on $[0, 1]$.*

The mapping $P(\cdot)$ is strictly decreasing. Given a higher frequency of disclosure, the model predicts a higher probability of a friction. Furthermore, because of strategic disclosure, the probability of the friction is always strictly less than the probability of not observing a disclosure.[7]

1.2 Intuition for the estimation strategy

We first explain the estimation strategy under the simplifying assumption that the probability of the friction is constant over time $p_t = p$. This assumption does not hold in practice and we shall later relax it by allowing for persistence in the friction. For now, we will assume that the market expectation μ_t is observed by the econometrician.

Suppose we have access to T observations for a given firm. The likelihood of the data is given by

$$\mathcal{L} = \begin{cases} 0 & \text{if } \frac{\min(\mathbf{x})}{\sigma} < Z(p) \\ \prod_t^T [g_t(v_{t+1}) \left(p + (1-p)\, H_t\left(Z(p)\right)\right]^{1-I_t} \left[(1-p)\, f_t(x_t, v_{t+1})\right)]^{I_t} & \text{if } \frac{\min(\mathbf{x})}{\sigma} \geq Z(p) \end{cases}.$$

where I_t is an indicator function which takes value 1 if there is disclosure in period t and zero otherwise. The function $g_t(v_{t+1})$ represents the marginal density of v_{t+1} in period t, and $H_t(.)$ is the distribution of x_t conditional on the realization of earnings

[7] We omit a formal discussion of identification but this can easily be shown in the iid case, for reasons that are analogous to those in Heckman (1979). For example, using the method of moments we see that Eq. (4) uniquely identifies the threshold z_t. Identification of σ^2 can then be obtained from $V_t(x_t|x_t > y_t) = \sigma^2(1 + z_t\delta(z_t) - (\delta(z_t))^2)$, where $\delta(x) = \frac{\phi(x)}{1-\Phi(x)}$. Finally, μ_t can be identified from $E[\mu_t|x_t > y_t] = \mu_t + \sigma\delta(z_t)$.

v_{t+1}. Since the model predicts that forecasts must be above the threshold y_t, the model enforces a lower bound for the dispersion of the manager beliefs

$$\sigma \geq \min(\mathbf{x})/Z(p). \tag{6}$$

This bound is key to understanding how DJK reconciles observations in which very bad news is voluntarily disclosed, that is, when $\min(\mathbf{x})$ is a large negative number. In the model, the *standardized* forecast must always be above $Z(p)$. So, for bad news to be disclosed, it must be that the manager signal is precise, because this leads to great variability in his posterior assessment. Intuitively, if the manager's information is precise, then the market will tend to penalize more a non-disclosure, and as a consequence, more unfavorable information will tend to be disclosed.

Naturally, a greater probability of friction p implies that the manager, overall, discloses fewer unfavorable news. Hence, as we have stated in claim 1, the increase in this bound is much greater when the frequency of disclosure is itself very low, because the low frequency indicated that this manager was unlikely to be informed. In summary, an occasional disclosure of bad news indicates a manager who is often subject to a friction but, when not, has very precise information to forecast.

The maximum likelihood estimator (p^{MLE}, σ^{MLE}) takes one of two forms. If equation (6) is not binding, the likelihood function can be decomposed to estimate σ^{MLE} and p^{MLE} separately, as

$$\begin{aligned}
\sigma^{MLE} &= \sqrt{\frac{\sum_t^T I_t(x_t - \mu_t)^2}{n}}, \\
p^{MLE} &= \arg\max_p \left(p + (1-p)\, H_t\left(Z(p)\right)\right)^{T-n} (1-p)^n, \\
\frac{\min(\mathbf{x})}{\sigma^{MLE}} &\geq Z(p^{MLE}),
\end{aligned}$$

where n is the number of periods where no disclosure was issued.

Note that, while the observed forecasts are truncated, the MLE estimator σ^{MLE} seemingly ignores the truncation and uses the estimator of an untruncated normal, even though a truncated normal has lower variance than a non-truncated one. The reason for this is tied to the equation for the bound in DJK: the disclosure threshold scales with the variance, that is, $y_t = \mu_t + \sigma Z(p)$.

Proper consideration of the bound is, however, critical when observing low forecasts. Suppose next that equation (6) binds, then the MLE estimator is

$$\begin{aligned}
\sigma^{MLE} &= \frac{\min(\mathbf{x})}{Z\left(p^{MLE}\right)}, \\
p^{MLE} &= \arg\max_{p} \prod_{t}^{T} g_t(v_{t+1}) \Big[p + (1-p)\, H_t(z)\Big]^{1-I_t} \Big[(1-p)\, f_t(x_t|v_{t+1})\Big]^{I_t}.
\end{aligned}$$

When equation (6) is binding, the estimation can no longer be conducted as two separate problems for p^{MLE} and σ^{MLE}. This is because to rationalize a history of infrequent but low forecasts we must either decrease the manager's precision or increase the probability of the friction p. The maximum likelihood estimator meets the bound by balancing this trade-off.

1.3 Stickiness in forecast policies

In practice, management forecast policies are sticky: the probability that a manager releases a forecast in the current quarter depends on how often he has released forecasts in the past. In their survey of management, Graham, Harvey and Rajgopal 2005 note that the act of disclosing information sets a precedent before investors whereby investors tend to expect more disclosure from a firm that has frequently disclosed in the past versus one that has remained silent.

To capture stickiness in disclosure decisions, we augment the model by assuming that the friction is persistent. Formally, assume that there is a process $\theta_t \in \{0, 1\}$ where $\theta_t = 1$ indicates that the manager is subject to the friction in period t. Specifically, we assume that θ_t follows a Markov chain with transition matrix:

$$K = \begin{bmatrix} k_1 & 1-k_1 \\ k_2 & 1-k_2 \end{bmatrix}.$$

That is, if (not) subject to the friction at date t, the manager is subject to the friction at date $t+1$ with probability (k_2) k_1. As before, investors do not observe the friction. Conditional on no-disclosure at date t, investors do not know whether $\theta_t = 1$ or, instead, $\theta_t = 0$ but the manager withheld his private information.

Investors update their beliefs about θ_t using the history of disclosure and their conjectures about the manager's equilibrium strategy. Conditional on observing a disclosure at date t, investors know that $\theta_t = 0$ (no friction), and update their belief for the next period based on the transition matrix, thus setting $p_{t+1} = k_2$. By contrast, conditional on observing no disclosure in period t, investors must update their beliefs to

$$\mathbb{E}_t(\theta_t | d_t = ND) = \underbrace{\frac{p_t}{p_t + (1-p_t)\mathbb{P}(x_t < Z(p_t)|v_{t+1})}}_{=\alpha(p_t)}.$$

Notice that $\alpha(p_t)$ uses the information the market observes in between disclosure periods, such as the earnings announcement, which naturally leads investors to reassess the probability that a manager withheld information. As an example, if realized earnings are low following a non-disclosure, investors will increase their posterior belief of strategic withholding implying, if the friction is sticky, a greater probability that the friction is not present in the next period.

We denote $d_t \in \{x_t, ND\}$ where $d_t = x_t$ indicates a forecast and $d_t = ND$ indicates no-disclosure. Then, the probability p_{t+1} follows:

$$p_{t+1} = \begin{cases} k_2 & \text{if} \quad d_t \neq ND \\ \alpha(p_t)k_1 + (1-\alpha(p_t))k_2 & \text{if} \quad d_t = ND \end{cases}. \tag{7}$$

This equation recursively characterizes the evolution of investors' beliefs. For the first observation in the sample, we use the steady state distribution of θ_t, namely we assume that $p_1 = \frac{k_2}{k_2+1-k_1}$. This approximation only affects the assessment prior to the first observed forecast because investors' beliefs are reset to k_2 each time a disclosure is observed.[8] As a special case, if $k_1 = k_2 = k$, the probability of receiving information is i.i.d., and the estimation procedure can be run as described in the previous section with a single state with $p = k_1 = k_2$.

We can now write the likelihood function as:

$$\mathcal{L}(K, \Sigma) = \prod_{t=1}^{T} [g_t(v_{t+1}) (p_t + (1 - p_t) H_t (Z(p))]^{1-I_t} [(1 - p_t)) f_t(x_t, v_{t+1}))]^{I_t}$$

where p_t follows the process in (7). Besides the variance covariance matrix, Σ, we must estimate the transition matrix of this Markov chain, K. The latter has a direct interpretation in terms of the stickiness of the underlying friction. Stickiness is measured by $k_2 - k_1$; when $k_2 = k_1$ the disclosure friction is *i.i.d.*

Based on the estimates of (K, Σ) we can compute two moments of special interest. First, we compute the steady-state probability of being subject to the friction which we denote by p_∞. That is,

$$(p_\infty, 1 - p_\infty) = (p_\infty, 1 - p_\infty)K.$$

This probability measures the importance of the friction and lies in between k_1 and k_2. This number tells us how often the manager does not disclosure for reasons unrelated to his information about the firm's future earnings.

Second, we define V_∞ as the steady-state probability that the manager withholds information for strategic reasons. Computing this parameter is more difficult than p_∞

[8]We have chosen this approximation because it is computationally tractable. An alternative would be to obtain by simulation the true likelihood function of the observations prior to the first forecast in steady-state; however, this method is computationally infeasible, however, given that we heavily rely on parametric bootstrap. When estimating the model on simulated data, the steady-state approximation does not seem to play a major role.

because it is an expectation many possible histories. We compute it by Monte-Carlo integration simulating histories of 1000 periods based on our MLE estimates, and recovering the probability as the average frequency of strategic withholding realized in the simulated histories.

It is well-known that, under various regularity conditions, the MLE estimator is consistent, asymptotically normal, and attains the Cramér-Rao lower bound (Amemiya, 1985). However, using the asymptotic properties of the estimator is not ideal in our setting because we have a small sample size of about 54 correlated observations per firm. So, we estimate standard errors by parametric bootstrap. Using the assumed parametric structure, we simulate data and generate a distribution of parameter estimates by assuming the firms are in steady state. Specifically, after estimating (K^{MLE}, Σ^{MLE}), we simulate 120 random histories of 500 observations, and re-estimate the model using the last 54 observations (because the original data has 54 quarters).

2 Estimation

2.1 Data

We focus our analysis on the set of firms who were in the S&P 500 index as of January 1st, 1998 which after several data restrictions reduces to a sample size of 371 firms. These firms are by construction large, with large analyst coverage and high institutional ownership and, consequently, unlikely to have missing observations from commercial vendors such as First Call (Chuk, Matsumoto and Miller, 2013).

Panel A of Table 1 describes the data filtering steps used in processing Firs Call's management forecasts. We begin with 22,617 forecasts from 456 firms. Next, we restrict analysis to fiscal years beginning in 1998 because of expanded coverage of the database to reduce selection bias (21,513 remaining forecasts) (Anilowski, Feng and Skinner 2007). We select only earnings per share forecasts (EPS) made prior to the end of the fiscal period to distinguish from pre-announcements (8,485 remaining forecasts). We select the first forecast if there was more than one. In most cases, this forecast comes bundled with the prior-quarter earnings release. Finally, we restrict forecasts to quantitative point and range estimates, leaving us with a sample of 4,750 forecasts from 371 firms. For range forecasts, we use First Call's normalized point estimate.[9]

Panel B of Table 1 describes the raw sample of earnings announcements, also from *First Call*.[10] We follow similar restrictions to the management forecasts - with the last step being a filter on firms that have at least 20 earnings announcement and 1 voluntary forecast in our entire sample period.[11] This final step results in a final sample of 371 firms making 4,750 earnings forecasts prior to 17,828 earnings announcements.

Panel A of Table 2 provides firm level summary statistics over the entire sampling period. The average firm experienced 54 reporting periods, with an average forecast frequency ($\hat{q}$) of 26%. The average realized earnings per share (split-adjusted) is .511, slightly greater than the average realized earnings conditional on no forecast .49. Also note that the mean forecast .59 is generally greater than the actual earnings .511, suggesting that the distribution of forecasts might be truncated. Panel B of Table 2 provides industry-means of key summary statistics (first and second moments of realized earnings and forecasts, as well as the average pre-forecast analyst consensus) where industry is defined using the 2-digit Global Industry Classification Scheme.

Two empirical choices are of importance, and worth discussing in more details. We use EPS, which is itself an issue of some debate in capital market research. Our approach

[9]When converting ranges to point estimates, First Call uses the mid-point of the range. For open-ended ranges (which are infrequent), *First Call* uses the bound plus one cent for upper ranges and minus one sent for lower ranges. Ranges did not appear to be extremely informative. In our sample, actual earnings fall in the stated range only about about one third of the time and the midpoint does not appear biased upward or downward.

[10]One benefit of using *First Call* is that the database normalizes its forecasts and realized earnings under the same measurement basis, making them comparable to analyst forecast data. These numbers are pro forma, that is, do not necessarily coincide with reported GAAP earnings in the Compustat database (see Barth *et al.* (2012) for more details). All earnings numbers used here are, as much as possible, under the same pro-forma basis.

[11]By construction we are excluding firms which almost never issue a forecast- though these firms are still rationalize-able under Dye's model e.g. as p approaches 1.

would be severely affected by market noise if we were, say, dividing by share price and, to be fair to the institutional setting, managers forecast EPS, not the price to earnings ratio. Cheong and Thomas (2011) show that earnings per share do not appear to show any variation with scale because, plausibly, firms manage their share price to ensure comparability. This is our implicit assumption when using earnings per share or, more accurately, the precision of the manager's private information about earnings per share does not significantly vary during the sample period.

We also checked whether managers bias their forecasts, that is, exhibit systematic non-zero forecast errors. There is some debate in capital market research as to whether management forecasts are biased; for example, Coller and Yohn (1998) conclude that forecasts do not appear biased while, with more recent data, Rogers and Stocken (2005) observe that forecasts appear biased. To address this, our estimation procedure includes a possible bias term. We do not find evidence of a significant positive or negative average bias, primarily because we focus on large firms. On a per-firm basis, evidence of a bias is mixed, with about one third of the firms with no significant bias, one third positive, and one third negative (Panel A Table 2).

2.2 Findings

To gain some flexibility in the estimation, we allow for the possibility that both the consensus μ_t and the management forecasts x_t are biased relative to the true earnings v_{t+1}. We define bias in forecast as $b_f \equiv E_t[x_t] - E_t[v_{t+1}]$ and bias in consensus as $b_\mu \equiv \mu_t - E_t[v_{t+1}]$.

Table 1 reports the estimated parameters and their standard errors, averaged across the sample firms. The friction appears to be fairly sticky, with 76% probability (k_1) of an existing friction persisting to the next quarter, versus a probability (k_2) of 52% of the friction realizing in the next quarter if it was not present in the current quarter.

Table 1. Average parameter estimates (N=371)

k_1	k_2	σ	φ	b_f	b_μ	Log-Likelihood
0.7567	0.5176	0.3307	0.6398	-0.1322	0.0220	1.9496
(0.1294)	(0.1577)	(0.1524)	(0.0829)	(11.7825)	(0.0833)	

To test the hypothesis of stickiness, we use a a chi-square test for nested models, taking $k_1 = k_2$ as the null hypothesis. When applying this test on a per-firm basis we find that the assumption of *i.i.d.* friction is rejected for 55% of sample firms (at 5% significance level). With regard to the quality of manager's information, we find that $\sigma^2 = 0.1089$, whereas the variation in earnings is $\varphi^2 = 0.41$. This implies that managers possess significant private information about the firm's prospects, but this information accounts only for 27% of the variation in earnings.

Is there evidence of strategic withholding? To answer this question, recall that under DJK, management forecasts are drawn from a truncated distribution. A truncated distribution has lower variance than the underlying non truncated distribution implying that, to explain the variance of management forecasts, the estimated variance of the manager's expectation must be greater than that of forecasts, with the variance gap increasing in the truncation point. This feature is consistent with some descriptive statistics. The variance of management forecasts is .0729 whereas the variance of manager expectations is $\sigma^2 = 0.1089$. A truncation in the distribution implies also that the average forecast should be larger than the average earnings, which again is supported by the data (.57 versus .51).

Table 2. Probability of Friction/Withholding (N=371)

p_∞	V_∞
0.6326	0.0951
(0.1055)	(0.0341)

Table 2 reports the steady-state probability of the friction and strategic withholding. The friction appears to be relatively frequent, with an average point estimate of 63%.

Also, we find that strategic withholding occurs only about 10% of the quarters, on average. The reason for this is tied to the assumptions of DJK: in this model market pressure forces disclosure of any above-average information so that, in the absence of the friction, strategic withholding must occur less than half of the time. Market pressure are indeed strong enough to induce disclosure of somewhat unfavorable events. Roughly, our estimates suggest that the manager conceals his information when it falls in the first quartile of the distribution of the manager's expectations (namely conditional on not facing the friction, the probability of withholding is $\frac{.0951}{1-.63}a \approx 25\%$.).

At first sight management forecasts appear to be downward biased, which echoes the idea that managers seek to influence the beliefs of analysts downwards in order to subsequently beat the consensus forecast via the actual earnings (Richardson, Teoh and Wysocki 2004). By contrast, the consensus forecast appears to be slightly upward biased on average, suggesting that analysts overstate their forecasts (for example Chen and Jiang (2006) conjecture that this is due to analyst incentives). However, we must point out that neither type of bias is statistically significant on average.

In Panel B, we break down these results by industry groups (2-digit GICS). We find that volatile industries or industries that rely on inputs that are hard to predict exhibit lower quality of managerial information, with energy and telecommunications featuring very low quality private information. On the other hand, as intuitively expected, utilities and industrial sectors feature greater quality. The case of the financial industry is somewhat interesting as well because (with telecommunications), it is one of the few industries for which we do not observe stickiness. The financial industry exhibits fairly precise managerial information, which may be driven by the role of accrual assumptions (in loan-loss provisions) quarter over quarter.

As to the steady-state probability of withholding, we find some moderate differences across industries. Telecommunications and utilities are the most likely to withhold information, about one quarter out of 5. By contrast, energy, healthcare and technology withhold information only about one quarter out of 10.

The last column presents the steady-state probability of experiencing a friction. The energy sector appears, by far, the most likely to exhibit a friction (79%), and information technology the least likely (27%). Our best conjecture for this result is that technology firms are the most likely to face short-term market pressure, as a result of stock compensation or their reliance on equity financing. But, then, we should also expect financial firms to be unlikely to have this friction; yet this is not the case (57%).

Figure 3 presents histograms for the parameter estimates across firms. Notice that the standardized disclosure threshold z is by construction negative, reflecting the prediction of DJK that only below-average news may be withheld. The histogram reveals a fair amount of variation at the firm level, especially for the disclosure threshold and the steady-state probability of the friction. Of practical interest, the model predicts a fairly small range for the estimated probability of withholding.

We note two additional facts. First, closer inspection of the threshold z reveals that a large peak at a disclosure around -0.3, that is, (by and large) most firms tend to strategically withhold information that is about one standard deviation below mean. Second, the histogram of σ^2 shows that a fraction of firms appear to have very low quality information (the mode of the distribution), essentially making a forecast that is unchanging - and therefore, which carries very little information.

3 Applications

3.1 A test of overconfidence

An extensive literature suggests that managerial over-confidence is widespread. The literature has used the term overconfidence in two senses: upward bias in beliefs (Malmendier, Tate and Yan, 2011) and over-sensitivity to private information (Chen and Jiang (2006); Marinovic, Ottaviani and Sørensen (2013)). Next, we study both phenomena but use the term over-confidence in the latter sense, namely excessive reliance on one's signal, whether positive or negative.

In general, disentangling overconfidence from private information is difficult. But the setting of management forecasts is well-suited for detecting over-confidence because manager forecasts are a direct report the manager's beliefs regarding a well defined event that is observed ex post, when earnings are revealed. To study over-confidence we rely on a standard test of excess volatility. Bayes' rule imposes significant structure on the joint distribution of forecasts and earnings realizations which one can exploit. For example, absent overconfidence, basic probability theory implies that $Var_t(x_t) = Var_t(\mathbb{E}_t(v_{t+1})) = \sigma^2 < Var_t(v_{t+1}) = \varphi^2$. That is, the variability in management forecasts must be lower than the variability in earnings.[12] A violation of this inequality can be explained under the assumption that the manager is overconfident and puts excessive weight on privately-informed information.

In order to formally test for over or under-confidence, we generalize the estimation DJK by estimating equation (2), without the rationality restriction $\rho = \sigma/\varphi$. A management forecast x_t can be written as

$$x_t = \phi_t + \delta \cdot \hat{x}_t, \tag{8}$$

where ϕ_t is predetermined, $\hat{x}_t$ represents the Bayesian beliefs conditional on the manager's private information and $\delta = \frac{\sigma}{\rho\varphi}$ represents overconfidence when $\delta > 1$ or underconfidence when $\delta < 1$.[13]

We label this generalized model as DJKO. The outcome of the estimation is presented in Table 3.

[12] Note that DJK remains applicable even with over-confidence, given that the input of the model is the manager's subjective beliefs about how the market react.

[13] As an equivalent formulation, assume the manager observes a signal $s_t = v_{t+1} + n_t$ where n_t is white noise. The manager forecast can then be written as $x_t = (1 - \delta\beta)\mu_t + \delta\beta(s_t - \mu_t)$ where $\beta = cov_t(s_t, v_{t+1})/var_t(s_t)$. Management forecasts are thus Bayesian if and only if $\delta = 1$. It would be possible to place additional structure and impose $\delta\beta \in [0, 1]$ but our estimates do not impose this restriction.

Table 3. Average parameter estimates (N=388)

	k_1	k_2	σ	φ	ρ	b_f	b_μ	Log-Lik
DJK	0.7567	0.5176	0.3307	0.6398	0.6011	-0.1322	0.0220	1.9496
DJKO	0.7756	0.5234	0.3816	0.5363	0.2967	-0.1338	-0.0158	19.5474

Notice the large increase in likelihood obtained upon allowing for over-confidence. We can estimate the average overconfidence parameter δ using the equality $\delta = \frac{\sigma}{\rho\varphi}$. The individual estimates of δ are, however, imprecise. 65% of the firms have a $\delta > 1$. The median is $\delta = 1.28$. The percentile 25% is 0.99 and the percentile 75% is 4.21. So, the evidence suggests that managers do not use their private information efficiently when they release a forecast and that, on average, they overreact to their private information.

Tests. To formally test for overconfidence, we use a likelihood ratio test, taking the DJK model as the null hypothesis H_0 and the overconfidence model $DJKO$ as the alternative hypothesis H_a. The test statistic is thus $\Lambda = \frac{\mathcal{L}(K^{MLE}, \Sigma^{MLE})}{\max_{K,\Sigma} \mathcal{L}_a(K,\Sigma)}$. Since our sample size per firm is relatively small, we use parametric bootstrap to compute the distribution of the test statistic using 120 draws. We also conduct a global test, where H_0: "DJK holds for all firms" against H_a "Model $DJKO$ holds for all firms." The test procedure is identical to the per firm test, except that we use the joint test statistic $\prod_i^N \Lambda_i$.

Table 4. Rejection of DJK vs DJKO

Reject DJK	Count	Percent
No	159	40.98%
Yes	229	59.02%

Table 5. Global test of DJK vs DJKO

Sig.	ΔLog-Lik	Crit.Value
1%	-17.599	-1.919
5%	-17.599	-1.610
10%	-17.599	-1.405

The results of the formal test are documented in Table 4. The test rejects rationality for as many as 59% of the firms. To further investigate the issue of overconfidence we conduct a global test of managerial rationality. Table 5 presents the results of a test that examines whether DJK can be applied to the entire universe of firms against DJKO which allows for over-confidence. Not surprisingly, at this point, rationality is rejected at any conventional level of significance.

3.2 What is the friction?

In this section, we develop an exploratory analysis of the economic forces that may drive the friction and the probability of strategic withholding. As we have seen, the friction has many possible interpretations which may cause the absence of disclosure.

The conventional explanation for the friction, given by Dye (1985), is that managers do not always receive incremental information - relative to investors - or receive it too late within the quarter. In Table 4, we observe weak evidence that firms with a greater likelihood of the friction p_∞ tend to feature less precise information (lower σ). Under the *assumption* that the probability of information endowment and the quality of the information are related metrics for the quality of managerial information, this correlation may suggest that a "friction" represents a situation where the manager does not have enough information to make a forecast.

Another explanation is that the friction may represent time-varying proprietary costs, in which making a forecast may occasionally cause large proprietary costs. In Table 4, we observe that firms that are larger - whose proprietary costs are likely to be greater - tend to be more likely to be subject to the friction. Indeed, Table 3 Panel B suggests that industries that are widely viewed to be less competitive, such as Energy, Utilities and Telecommunications, are more likely to be subject to the friction.

An inspection of the time-series of our estimation reveals another likely determinant of the friction. In Figure 2, we plot the average beliefs about the friction as implied by the estimated DJK across fiscal reporting periods for the average firm. We observe a

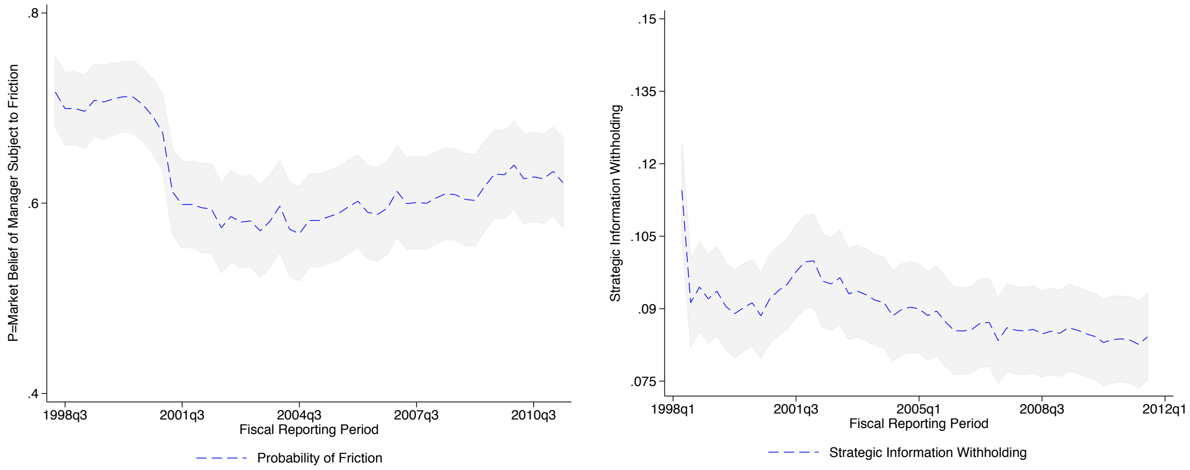

Figure 2. Evolution Market Beliefs and Probability of Strategic Withholding. This figure plots the average probability of the market's belief that the manager is subject to the friction averaged across firms. The gray regions represent 95% confidence intervals.

very clear downward shock affecting the probability of the friction between $2000 - q3$ and $2001 - q4$. which coincides with greater stock market pressure because of the large corporate frauds that were unveiled (e.g., Enron and WorldCom), which ultimately led to the enactment of the Sarbanes Oxley act in 2002. It also coincides with Regulation Fair Disclosure, a regulation that was promulgated by the U.S. Securities and Exchange Commission in August 2000, mandating that all publicly traded companies must disclose material information to all investors at the same time, rather than selectively to some analysts. The regulation prevented firms from reporting privately to analysts, causing greater investor pressure for public management forecasts as a substitute.

4 Conclusion

In this study, we develop a simple methodology to estimate the disclosure model of Dye (1985) and Jung and Kwon (1988) using management forecast data.

This is the first study to estimate the probability of strategic withholding of information in capital markets and to separate the strategic and non-strategic components of non disclosure. We have documented the following facts. 1) Managers strategically withhold bad news about 10% of the time. 2) They conceal the percentile 25% of the distribution of private information in the absence of the friction. 3) They enjoy a significant informational advantage, relative to the market: their private information at the disclosure date explains 50% of the variation in earnings. 4) We provide evidence that managers on average over-react to their private information.

We hope to offer, here, a starting point to identify the failures of the disclosure theory and offer as much material for further research in empirical and theoretical research. The data appears to suggest that there are fundamental differences between how firms forecasts, with some following DJK while others disclosing information regardless of its price impact. Further work should, therefore, involve richer settings that improve over the DJK approach. For theoretical research, we appear to find robust evidence that managers do not form rational expectations but, instead, feature some large levels of overconfidence. Lastly, our model is amenable to test richer theories to disclosure by conditioning the probability of the friction on various determinants of disclosure.

References

Acharya, V. V., DeMarzo, P. and Kremer, I. (2011). Endogenous information flows and the clustering of announcements. *The American Economic Review*, **101** (7), 2955–2979.

— and Lambrecht, B. M. (2011). *A theory of income smoothing when insiders know more than outsiders*. Tech. rep., National Bureau of Economic Research.

Amemiya, T. (1984). Tobit models: A survey. *Journal of econometrics*, **24** (1), 3–61.

— (1985). *Advanced econometrics*. Harvard university press.

Anilowski, C., Feng, M. and Skinner, D. J. (2007). Does earnings guidance affect market returns? The nature and information content of aggregate earnings guidance. *Journal of Accounting and Economics*, **44**, 36–63.

Bajari, P. and Hortacsu, A. (2003). The winner's curse, reserve prices, and endogenous entry: empirical insights from ebay auctions. *RAND Journal of Economics*, pp. 329–355.

Ball, R. and Brown, P. (1968). An empirical evaluation of accounting income numbers. *Journal of accounting research*, pp. 159–178.

Barth, M. E., Gow, I. D. and Taylor, D. J. (2012). Why do pro forma and street earnings not reflect changes in gaap? evidence from sfas 123r. *Review of Accounting Studies*, **17** (3), 526–562.

Ben-Porath, E., Dekel, E. and Lipman, B. L. (2014). *Disclosure and Choice*. Tech. rep., Mimeo.

Beyer, A., Cohen, D. A., Lys, T. Z. and Walther, B. R. (2010). The financial reporting environment: Review of the recent literature. *Journal of Accounting and Economics*, **50** (2-3), 296–343.

—, Guttman, I. and Marinovic, I. (2014). Earnings management and earnings quality: Theory and evidence. *Available at SSRN 2516538*.

Bushman, R. M. (1991). Public disclosure and the structure of private information markets. *Journal of Accounting Research*, pp. 261–276.

Chen, Q. and Jiang, W. (2006). AnalystsŠ weighting of private and public information. *Review of Financial Studies*, **19** (1), 319–355.

Cheong, F. S. and Thomas, J. (2011). Why do eps forecast error and dispersion not vary with scale? implications for analyst and managerial behavior. *Journal of Accounting Research*, **49** (2), 359–401.

Chuk, E., Matsumoto, D. and Miller, G. S. (2013). Assessing methods of identifying management forecasts: CIG vs. researcher collected. *Journal of Accounting and Economics*, **55** (1), 23–42.

Coller, M. and Yohn, T. L. (1998). Management forecasts: What do we know? *Financial Analysts Journal*, **54** (1), 58–62.

Dickhaut, J., Ledyard, M., Mukherji, A. and Sapra, H. (2003). Information management and valuation: an experimental investigation. *Games and Economic Behavior*, **44** (1), 26–53.

Dye, R. A. (1985). Disclosure of nonproprietary information. *Journal of Accounting Research*, **23** (1), 123–145.

— (1988). Earnings management in an overlapping generations model. *Journal of Accounting Research*, **26** (2), 195–235.

— and Sridhar, S. S. (1995). Industry-wide disclosure dynamics. *Journal of accounting research*, **33** (1), 157–174.

Erdem, T. and Keane, M. P. (1996). Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing science*, **15** (1), 1–20.

Fischer, P. E. and Verrecchia, R. E. (2000). Reporting bias. *The Accounting Review*, **75** (2), 229–245.

Fishman, M. J. and Hagerty, K. M. (1990). The optimal amount of discretion to allow in disclosure. *Quarterly Journal of Economics*, **105** (2), 427–444.

Ganuza, J.-J. and Penalva, J. S. (2010). Signal orderings based on dispersion and the supply of private information in auctions. *Econometrica*, **78** (3), 1007–1030.

Gentzkow, M. and Shapiro, J. M. (2010). What drives media slant? evidence from us daily newspapers. *Econometrica*, **78** (1), 35–71.

Goldman, E. and Slezak, S. L. (2006). An equilibrium model of incentive contracts in the presence of information manipulation. *Journal of Financial Economics*, **80** (3), 603–626.

Govindaraj, S., Lee, P. and Tinkelman, D. (2007). Using the event study methodology to measure the social costs of litigation-a re-examination using cases from the automobile industry. *Review of Law & Economics*, **3** (2), 341–382.

Graham, J. R., Harvey, C. R. and Rajgopal, S. (2005). The economic implications of corporate financial reporting. *Journal of Accounting and Economics*, **40** (1), 3–73.

Grossman, S. J. (1981). The informational role of warranties and private disclosure about product quality. *Journal of Law and Economics*, **24** (3), 461–483.

— and Hart, O. D. (1980). Disclosure laws and takeover bids. *Journal of Finance*, **35** (2), 323–334.

Guttman, I., Kremer, I. and Skrzypacz, A. (2014). Not only what but also when: A theory of dynamic voluntary disclosure. *The American Economic Review*, **104** (8), 2400–2420.

Hagerty, K. M. and Fishman, M. J. (1989). Disclosure decisions by firms and the competition for price efficiency. *Journal of Finance*, **44** (3), 633–646.

Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica: Journal of the econometric society*, pp. 153–161.

Huddart, S., Hughes, J. S. and Levine, C. B. (2001). Public disclosure and dissimulation of insider trades. *Econometrica*, **69** (3), 665–681.

Jin, G. Z., Luca, M. and Martin, D. (2015). Is no news (perceived as) bad news? an experimental investigation of information disclosure. *An Experimental Investigation of Information Disclosure (March 31, 2015). Harvard Business School NOM Unit Working Paper*, (15-078).

Jovanovic, B. (1982). Truthful disclosure of information. *Bell Journal of Economics*, **13** (1), 36–44.

Jung, W.-O. and Kwon, Y. K. (1988). Disclosure when the market is unsure of information endowment of managers. *Journal of Accounting Research*, **26** (1), 146–153.

Kartik, N., Ottaviani, M. and Squintani, F. (2007). Credulity, lies, and costly talk. *Journal of Economic theory*, **134** (1), 93–116.

Kedia, S. and Philippon, T. (2009). The economics of fraudulent accounting. *Review of Financial Studies*, **22** (6), 2169–2199.

LeRoy, S. F. and Porter, R. D. (1981). The present-value relation: Tests based on implied variance bounds. *Econometrica*, pp. 555–574.

Lundholm, R. J. and Rogo, R. (2014). Do analysts forecasts vary too much? *Available at SSRN.*

Malmendier, U., Tate, G. and Yan, J. (2011). Overconfidence and early-life experiences: the effect of managerial traits on corporate financial policies. *The Journal of finance*, **66** (5), 1687–1733.

Marinovic, I., Ottaviani, M. and Sørensen, P. N. (2013). ForecastersŠ objectives and strategies. *Handbook of economic forecasting*, **2**, 691–720.

McClellan, F. (2008). The vioxx litigation: A critical look at trial tactics, the tort system, and the roles of lawyers in mass tort litigation. *DePaul Law Review*, **57**, 509–538.

Mello, M. M., Rimm, E. B. and Studdert, D. M. (2003). The mclawsuit: The fast-food industry and legal accountability for obesity. *Health Affairs*, **22** (6), 207–216.

Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *Bell Journal of Economics*, **12** (2), 380–391.

Rabin, R. L. (2001). Tobacco litigation: A tentative assessment. *DePaul Law Review*, **51**, 331.

Richardson, S., Teoh, S. H. and Wysocki, P. D. (2004). The walk-down to beatable analyst forecasts: The role of equity issuance and insider trading incentives*. *Contemporary Accounting Research*, **21** (4), 885–924.

Rogers, J. L. and Stocken, P. C. (2005). Credibility of management forecasts. *The Accounting Review*, **80**.

Shavell, S. (1994). Acquisition and disclosure of information prior to sale. *Rand Journal of Economics*, **25** (1), 20–36.

Shiller, R. J. (1980). Do stock prices move too much to be justified by subsequent changes in dividends?

Shin, H. (2003). Disclosures and asset returns. *Econometrica*, **71** (1), 105–133.

Shin, H. S. (1994). News management and the value of firms. *Rand Journal of Economics*, pp. 58–71.

Stein, J. C. (1989). Efficient capital markets, inefficient firms: A model of myopic corporate behavior. *Quarterly Journal of Economics*, **104** (4), 655–669.

Stocken, P. C. (2000). Credibility of voluntary disclosure. *The RAND Journal of Economics*, pp. 359–374.

Terry, S. J. (2014). *The Macro Impact of Short-Termism.* Tech. rep., Working paper.

Tobin, J. (1958). Estimation of relationships for limited dependent variables. *Econometrica: journal of the Econometric Society*, pp. 24–36.

Zakolyukina, A. A. (2014). Measuring intentional gaap violations: A structural approach. *Chicago Booth Research Paper*, (13-45).

Table 1. Sample Selection and Summary Statistics

Our firm sample is the set of firms who were initially in the S&P 500 index as of January 1st, 1998 for the sample period between fiscal quarter 1 of 1998 to fiscal quarter 3 of 2011. Panel A and B report our sample selection criterion of management forecasts and earnings announcements from the First Call database.

Panel A: Management Forecasts

Step	Rule	Obs	Firms
	Raw sample	22,617	456
Restrict to:			
2)	1998Q1-2011Q3 fiscal years	21,513	
3)	Quarterly guidance on EPS	8,485	
4)	Guidance date prior to end of FP	7,167	
5)	Guidance date within 90 days window prior to end of FP	6,508	
6)	Non-missing First Call point estimate	5,944	
7)	First dated guidance per fiscal period	4,750	371

Panel B: Earnings Announcements

Step	Rule	Obs	Firms
	Raw sample	54,775	495
Restrict to:			
1)	1998Q1-2011Q3 fiscal years	32,902	
2)	Original un-restated announcements	19,148	
4)	Firms with minimum of 1 voluntary forecast	17,828	371

Table 2. Summary Statistics of Raw Data

Panel A reports key summary statistics of our final sample of 371 firms. Management forecasts and realized earnings are both split-adjusted. The pre-forecast analyst consensus is defined as the average consensus (split-adjusted) of the most current (on an analyst by analyst basis) forecasts made prior to the date of management forecast disclosure. Panel B groups firms according to their historical 1998 two-digit Global Industry Classification Scheme (GICS) and reports the average of firm-level forecast frequencies, and their average realized earnings when preceded by disclosure or non-disclosure of a management forecast.

Panel A: Firm Level

Variable	Obs	Mean	Std. Dev.	P25	P50	P75
Number of reporting periods	371	46.77	13.13	43	54	54
Fraction of periods with forecasts ($\hat{q}$)	371	.26	.24	.06	.15	.44
Average value of realized earnings	371	.51	.54	.29	.46	.65
Average value of realized earnings (forecast disclosed)	371	.51	.65	.26	.42	.64
Average value of realized earnings (forecast not disclosed)	371	.49	.54	.25	.44	.64
Standard deviation of realized earnings	371	.5	1.29	.16	.26	.42
Average value of forecasts	371	.59	.62	.31	.5	.76
Standard deviation of forecasts	311	.27	.26	.11	.19	.36
Average mgmt. forecast error	371	.08	.74	-.02	0	.11
St.dev. of mgmt. forecast error	311	.23	.76	.04	.09	.21
Pre-forecast analyst consensus	371	.56	.63	.31	.48	.66

Panel B: Industry Level

GICS2	N	Avg frequency of forecasts	Avg realized earnings (forecast disclosed)	Avg realized earnings (forecast not disclosed)
Energy	11	0.14	0.52	0.66
Materials	36	0.22	0.47	0.43
Industrials	65	0.29	0.50	0.51
Consumer Discretionary	74	0.34	0.49	0.41
Consumer Staples	33	0.22	0.47	0.48
Health Care	28	0.31	0.57	0.47
Financials	37	0.07	0.66	0.69
Information Technology	44	0.41	0.44	0.41
Telecommunication Services	9	0.14	0.44	0.35
Utilities	25	0.11	0.57	0.57
Unclassified	9	0.32	0.44	0.44
Total	371	0.26	0.51	0.49

Table 3. Summary Statistics of Model Estimates – Firm Level

Panel A summarizes the key estimates of our three models (in superscripts, null, model "a" non-strategic, and model "b" naive-investors). k_1 (k_2) is the probability of the manager not bound by the disclosure friction conditional on being not bound (bound) in the previous period. σ is the standard deviation from the true (un-truncated) distribution of management forecasts. v_∞ is the steady-state probability that the manager is withholding news and p_∞ is the steady-state probability that the manager is bound by the disclosure friction. Panel B groups firms according to their historical 1998 two-digit Global Industry Classification Scheme (GICS) and reports the average of the key model estimates by these groupings (note that these are not model re-estimates at the industry level).

Panel A: Model Estimates – Firm Level

Variable	Obs	Mean	Std. Dev.	P25	P50	P75
k_1	371	.76	.25	.69	.86	.94
k_2	371	.52	.36	.16	.5	.91
σ	371	.33	2.45	.05	.11	.24
φ	371	.64	3.06	.1	.19	.36
p_∞	371	.63	.32	.35	.76	.91
v_∞	371	.09	.07	.03	.08	.15
Model Likelihood	371	1.95	61.28	-24.75	9.52	42.59
ρ	371	.6	.37	.52	.7	.82

Panel B: Model Estimates – Industry Level

GICS2	N	k1	k2	sigma	varphi	p	V	LIK	rho
Energy	11	0.82	0.56	0.24	0.31	0.77	0.08	-3.90	0.76
Materials	38	0.79	0.59	0.24	0.36	0.70	0.08	-9.04	0.64
Industrials	64	0.75	0.46	0.37	1.09	0.58	0.10	-15.36	0.64
Consumer Discretionary	76	0.69	0.37	0.83	1.07	0.50	0.12	-2.96	0.66
Consumer Staples	33	0.79	0.55	0.09	0.13	0.67	0.09	43.55	0.65
Health Care	29	0.76	0.52	0.11	0.16	0.65	0.12	35.25	0.66
Financials	40	0.86	0.79	0.21	0.76	0.84	0.05	-29.09	0.46
Information Technology	45	0.63	0.33	0.22	0.63	0.44	0.12	14.63	0.72
Telecommunication Services	10	0.91	0.58	0.07	0.32	0.80	0.06	29.55	0.44
Utilities	27	0.89	0.74	0.13	0.26	0.83	0.06	-0.27	0.44
Unclassified	15	0.68	0.61	0.04	0.14	0.68	0.08	7.28	0.12
Total	371	0.76	0.52	0.33	0.64	0.63	0.09	1.95	0.60

Table 4. Cross-Correlation of Estimated Parameters

This table reports the Pearson correlation among parameters estimates from Panel A of Table 3. k_1 (k_2) is the probability of the manager not bound by the disclosure friction conditional on being not bound (bound) in the previous period. σ is the standard deviation from the true (un-truncated) distribution of management forecasts. v_∞ is the steady-state probability that the manager is withholding news and p_∞ is the steady-state probability that the manager is bound by the friction. P-values are reported in parentheses below each correlation.

	log(Assets)	σ	φ	ρ	p_∞	V_∞	Log-Lik
log(Assets)	1						
σ	0.0821	1					
φ	0.0296	0.5553*	1				
ρ	-0.0235	0.2644*	-0.0781	1			
p_∞	0.1347*	-0.2063*	-0.0712	-0.4154*	1		
V_∞	-0.1388*	0.2793*	0.1525*	0.4658*	-0.6314*	1	
Log-Lik	-0.1374*	-0.5991*	-0.6032*	0.0361	0.0964	-0.1975*	1

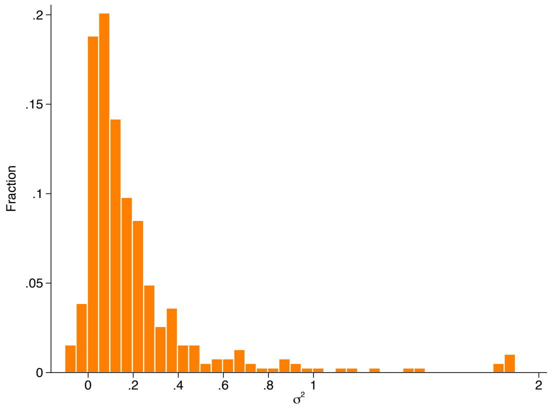

(a) $\hat{\sigma^2}$

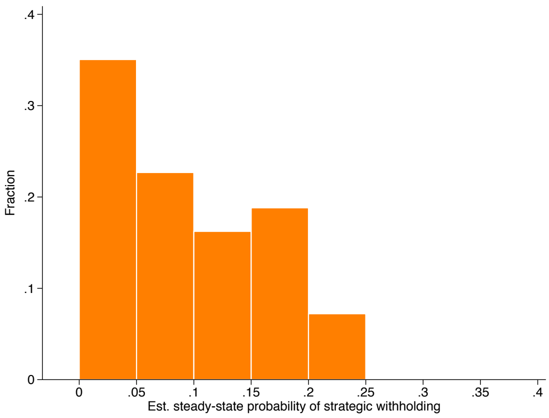

(b) Steady-state probability of strategic withholding

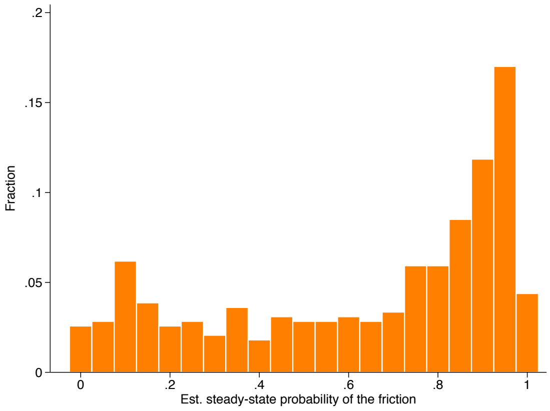

(c) Steady-state probability of the friction

Figure 3. Distribution of Estimated Parameters Under DJK. These figures plot the distribution of firm level estimates for $\hat{z}$, the disclosure threshold, $\hat{p}_{i,\infty}$, the steady-state probability that the manager is bound by the disclosure friction, $v_{i,\infty}$, the steady-state probability the manager is strategically withholding, and $\hat{\sigma^2}$, the estimated precision of the manager's information.