The Dynamics of Concealment

Jeremy Bertomeu, Iván Marinovic, Stephen J. Terry, and Felipe Varas *

November 8, 2019

Abstract

Management likely possess information about firm prospects that outside investors do not have. If managers strategically conceal such information, the accuracy of firm valuations may decline. To quantify this channel, we develop and solve a dynamic model of corporate disclosure. We estimate our structural model using a comprehensive sample of management earnings forecasts released by public companies in the US between 2004 and 2016. We estimate that about half of all firms strategically manage their disclosure. Strategic disclosers are endowed with information about 60% of the time and withhold information 26% of the time when informed, increasing uncertainty by roughly 4% of total earnings innovations relative to an environment with no strategic concealment. Information endowment is highly persistent, allowing investors to assess the existence of information in future periods after a first disclosure but also deterring managers from initiating a disclosure after a long non-disclosure spell. The model also explains why additional information from financial analysts encourages more forthcoming forecasts. Overall, our setting clarifies the benefits, but also inherent limits, of voluntary financial reporting.

Keywords: voluntary disclosure; structural estimation; reputations; persuasion.
JEL Classification: D72; D82; D83; G20.

*J. Bertomeu is at UC San Diego, I. Marinovic is at the Stanford Graduate School of Business, S. Terry is at Boston University, and F. Varas is at the Duke Fuqua School of Business. We thank Anne Beyer, Oeystein Daljord, Ilan Guttman, Paul Fischer, and Marco Ottaviani and seminar participants at Bocconi, Carnegie Mellon, Columbia (Burton workshop), and Tilburg for helpful comments.

1 Introduction

Where does news about future cash flows come from? In the US, public firms release their financial statements each quarter but accounting earnings are largely anticipated by investors before earnings announcements (Ball and Brown 1968). In this study, we investigate a voluntary channel used by corporations to transmit forward-looking information. Firms often supplement their financial statements with voluntary forecasts about future earnings. Empirically, these forecasts contain most of the price-relevant news released in recurrent reporting events. For example, Beyer, Cohen, Lys and Walther (2010) document that management forecasts explain a large portion of a firm's quarterly return variance, about 16%, more than analyst forecasts, earnings announcements, and regulatory filings combined.

Presumably, managers have information about their company that they may be willing to transmit before earnings announcements. But such voluntary forecasts also come with incentives to strategically withhold information. Evaluating the process through which information is received and selected for public disclosure is the main theme of this paper. We ask several questions. How much more information do managers possess that investors do not have, and what portion of this information is concealed? To what extent does strategic withholding cause deviations between market prices and fundamentals and what increases in price efficiency could we expect if strategic withholding were eliminated? Lastly, what factors contribute to make managers' disclosures more forthcoming?

A practical challenge to addressing these questions is that we do not directly observe the information received by management; instead, we observe the output of a disclosure process that selects which information is publicly reported. Therefore, we rely on a structural model linking observed disclosures to assumptions about the choices made by management and how the market forms its expectation after spells of non-disclosure. Furthermore, management forecasts are a dynamic process that depends on past news, being affected by past forecasts and realized earnings as well as analysts' information. In turn, this suggests benefits to using a dynamic model aimed at targeting facts of both of the cross-section and time-series of earnings, and in which forward-looking considerations by management play a key role.

Our theory is consistent with conventional stylized facts about the empirical distribution of forecasts and earnings. While unravelling theory predicts that all firms should make a forecast (Grossman and Hart 1980; Milgrom 1981), we observe a low propensity to forecast. Even for large firms with active investor relations, a forecast is made only between one third to one half of the time. We also observe a significant stickiness in propensities to disclose which, empirically, is poorly explained by observed firm characteristics. Managers appear to select periods in which information is more favorable to make a forecast and strategically withhold information (Kothari, Shu and Wysocki 2009). Markets react unfavorably to firms that stop giving forecasts and such decisions usually precede declines in actual earnings.

Before we lay out our theoretical framework, we further discipline our analysis with three empirical facts that our theory will attempt to quantitatively explain. For expositional purpose, we discuss these facts here in broad strokes, deferring a complete discussion within our sample until section 2.

(1) Disclosures are far from exogenous draws from an earnings process, in contrast to what would be expected if disclosures were mandatory disclosure events, e.g., an 8-K filing for material news. Realized earnings are almost 1% higher when they are preceded by a voluntary forecast, relative to earnings during long spells without forecasts. This fact suggests that managers strategically select which future earnings to forecast based on private information. This effect flattens around 3 years from a forecasted year and implying that there are dynamic effects linking current forecasts to past and future earnings.

(2) When disclosure policies change, especially in the case of a forecast following a long spell of non-disclosure or, vice-versa, a non-disclosure following a long spell of forecasts, there is a noticeable break in the earnings process, which carries over current and future period earnings. Empirically, current earnings decrease by about 12% of lagged earnings following a break in disclosure behavior, reducing the effect of lagged

earnings by about a fourth. Breaks in disclosure behavior convey information about fundamentals that partly subsumes lagged earnings.

(3) Disclosures are more sensitive to earnings when preceded by a history of disclosures. Each percent of asset increases the probability of a forecast by about one percent when it follows another forecast, versus one fourth of a percent when there is no forecast. This is a key unexplained fact in the empirical literature: namely, firms appear to dynamically form reputations not to disclose but the mechanisms sustaining such reputations are not fully understood. This fact suggests a dynamic model in which price responses endogenously sustain reputations.

We rationalize these facts as an optimal strategic disclosure policy with persistent information endowment. We enrich the model of Dye (1985) and Jung and Kwon (1988), a static model in which a manager may not be endowed with information and maximizes the current price. To make this model amenable to an empirical analysis, we make this model fully dynamic, adding to it insights from recent disclosure theory: (a) the manager is forward looking and considers the effect of a disclosure on future prices (Guttman, Kremer and Skrzypacz 2014; Marinovic, Skrzypacz and Varas 2015; Beyer and Dye 2012); (b) managers' information endowment is serially correlated, and markets update their beliefs as a function of realized earnings (Einhorn and Ziv 2008); and (c) a public news process, e.g. analyst consensus forecasts, affects the propensity to disclose (Acharya, Demarzo and Kremer 2011).

The model implies a threshold level of internal news about profitability below which an informed manager chooses to conceal information from the market. This threshold is time-varying and depends on the market assessment of the current probability that the manager is informed, which is itself a function of the entire past history of forecasts. A forward-looking manager benefits from maintaining a reputation to be uninformed because the market penalizes non-disclosure less when strategic withholding is perceived to be unlikely. By making a forecast, the manager reveals their information endowment in the current period, and, thus, a higher likelihood that they will be informed in future. The loss of reputation that follows in our model generates an endogenous cost of disclosure, especially after long periods without a forecast. Rationalizing the survey evidence from Graham, Harvey and Rajgopal (2005), managers are reluctant to initiate a forecast because doing so would create a precedent that would increase pressure to make forecasts in future. Therefore, after long periods without forecasts, the firm will be less likely to disclose and more likely to reveal only the most favorable information to the public.

Our approach provides a quantitative assessment of a channel through which mandatory disclosure of realized earnings affects the provision of voluntary disclosure. In a one-shot disclosure model, a public signal like realized earnings would not affect the probability of disclosure because disclosure strategies are invariant to the mean and variance of posterior beliefs (Acharya, Demarzo and Kremer 2011). In our full dynamic model, by contrast, even if actual earnings contain little news, they help discipline forecasting behavior. Actual realized earnings inform investors about the private information observed by the manager and reduce a manager's ability to build an uninformed reputation. When observing low realized earnings, the market knows that the manager was more likely to be informed and will carry this information into the future. The greater the information contained in earnings, the less likely it is that the manager will be able to maintain a misleading reputation.

We estimate our structural model using a comprehensive dataset of annual realized earnings, management forecasts, and analyst consensus forecasts provided by I/B/E/S. Our sample is a panel of $8,109$ public firms spanning $2004-2016$ for a total of $53,167$ firm-fiscal years. We estimate the volatility of fundamental earnings and the precision of information contained in analyst forecasts by directly matching moments of observed earnings and consensus forecasts. We also estimate the likelihood that a manager is informed, together with the persistence of manager information, by requiring in a straightforward simulated method of moments (SMM) procedure that the model closely reproduce the likelihood of manager disclosure, the persistence of disclosure policies, and the observed magnitude of manager forecast errors.

Our baseline point estimates indicate that managers with potential strategic motives possess more information than investors about 60% of the time and strategically conceal this information with probability 26%. Relative to a counterfactual in which managers report all their information, strategic withholding increases the root mean squared error (RMSE) of investor earnings expectations by 9% of the total earnings innovation. These estimates point to policymaker and academic concerns with strategic concealment and suggests that disclosure policies and incentives are critical for determining price efficiency in public markets.

We build on a substantial body of previous theoretical work on disclosure. Our dynamic model is a part of what Milgrom (1981) defines as persuasion theory, namely, a class of sender–receiver problems in which the sender's preference depends only on the receiver's posterior belief. For the most part, the early disclosure literature is static (Jovanovic 1982; Verrecchia 1983; Dye 1985; Shin 1994; Shavell 1994). Persuasion theory has only been recently extended to a dynamic context because most single-period disclosure models no longer fit this definition of persuasion when they are repeated.

Three recent examples of two-period disclosure models serve to best illustrate this point. Acharya, Demarzo and Kremer (2011) develop a model in which the manager can delay a disclosure after the release of public news. In this model, the value of delay is a function of the manager's expectation about the public news, which depends on the information privately observed by the manager. Beyer and Dye (2012) consider a model in which some managers may be forthcoming and disclose all of their information. Managers know their propensity to be forthcoming in the next period, which implies that their type affects their payoff from withholding. Guttman, Kremer and Skrzypacz (2014) examine a model in which more information may be received at some later date. A manager choosing to delay can anticipate the information received at a later date: managers with a low signal know it is more likely that they will be willing to withhold later. The traditional methods of persuasion theory fail to apply or, at least, are significantly changed for repeated versions of these models.

One approach has been shown to preserve the basic structure of persuasion with some dynamic aspects. Einhorn and Ziv (2008) and Marinovic (2012) are models in which the market updates dynamically to new information, but the manager is myopic and maximizes only current short-term stock prices. The stock price is the sole channel through which future periods affect current actions. In Einhorn and Ziv (2008), the stock price includes the discounted value of expected disclosures in the future. In Marinovic (2012), the stock price is equal to the value implied by all past disclosures. After observing a current disclosure, the market reassesses the propensity to be truthful and reprices the entire sequence of past disclosures.

Our model shares its focus with a recent literature applying structural models to analyze strategic financial communication. Bertomeu, Ma and Marinovic (2016b) is the paper closest to ours and estimates the Dye (1985) model under the assumption that the manager, when informed, maximizes current stock prices. This paper generalizes the approach to a manager with forward-looking motives and adds an exogenous public news process, analysts' forecasts, to the estimation procedure. Several antecedents to our paper estimate strategic withholding with disclosure costs rather than uncertainty about information endowment (Bertomeu, Beyer and Taylor 2016a; Zhou 2017; Cheynel and Liu 2017). However, their research question is quite different because the probability of strategic withholding must be equal to the empirical probability of non-disclosure - their focus on the estimation and implications of disclosure costs.

Several other studies focus on incentives to manipulate the earnings process rather than selectively withhold forecasts. Beyer, Guttman and Marinovic (2014) structurally estimate a dynamic model of costly earnings misreporting, in which the manager's equity incentives are not known to the market (Fischer and Verrecchia 2000; Frankel and Kartik 2014). Zakolyukina (2014) examines a model in which managers trade off higher stock prices in the present versus a greater probability of prosecution in the future when making manipulation or deception choices. Terry (2016) and Terry, Whited and Zakolyukina (2018) study models in which manager earnings misreporting choices interact with real manipulation of long-term investments. We share with these models a focus on measuring information loss due to managerial discretion over the information flow; however, these earlier studies examine manipulation noise in reported accounting numbers while we focus on the voluntary disclosure channel.

2 Three stylized Facts

We start our analysis by documenting three stylized facts about manager's forecasting behavior and its association with earnings and analyst forecasts. Our objective is two-fold: rely on explanations from theory to discipline our model assumptions and formalize which key aspects of the data we aim at matching in the structural model. The sample used to recover these facts consists of public firms over the period $2004 - 2016$, making voluntary disclosures in the form of annual earnings forecasts; all details pertaining to the sample construction and variable definitions are deferred to section 4. Note that we do not use any of these stylized facts directly as identifying restrictions in the estimation so that they may provide additional validation of the model.

Fact 1: Selective Forecasting

Disclosure theory predicts that firms follow upper-tail disclosure strategies, in which news that is not sufficiently favorable is concealed (Jovanovic 1982; Verrecchia 1983; Dye 1985; Jung and Kwon 1988). Specific to truthful communication with frictions, this prediction is quite different from other modes of communication, such as cheap talk or costly signalling, in which there is no clear prediction about periods of non-disclosure or selection over which information to report.

Empirically, we observe that very few firms apply a policy of consistently forecasting all periods. For example, 90% of all firms choose not to issue annual forecasts in certain years and the average frequency of disclosure for firms having disclosed at least once is 59%.

While this fact alone does not imply that firms strategically conceal bad news, we show in Figure 1 how forecast disclosures coincide with higher realized earnings for the forecasted year. For each firm-year, we take the difference between current earnings and earnings three years earlier conditional on the existence of a forecast (d_t). As implied by the theory, realized earnings are higher, by about .75% of current assets.

We repeat this exercise conditional on disclosures after the forecasted year (d_{t+1}, d_{t+2} , d_{t+3}) and before the forecasted year ($d_{t-1}, d_{t-2}, d_{t-3}$), with a 95% confidence intervals indicated as a dotted line. Because earnings are serially-correlated, we expect higher earnings to peak at the earnings date and decline as disclosures become more distant. This is also what we observe empirically, with the peak in earnings declining by about half one year from the forecasted year. At about 3 years ahead from the forecasted year, current earnings are about equal to what they were three years before the forecast. This fact motivates a model with strategic reporting as, absent strategic motives, the empirical relationship between disclosure and earnings would be flat.

Fact 2: Effect of breaks in disclosure policy

The previous fact suggests that firms may strategically withhold bad news, but it does not tell us anything about managers' forward looking behavior or the long run consequences of withholding a forecast.

Sometimes, a firm breaks its past disclosure policy, by for instance releasing a forecast when it has not done it in the past, or refraining from doing so when it did it in the previous period. Such breaks contain information about the upcoming earnings. To document this fact, we construct an indicator variable indicating a break in disclosure policy, equal to one when a firm switches from disclosure to non-disclosure (*dend*) or, vice-versa, from non-disclosure to disclosure (*dstart*). Then, we estimate the following regression:

$$\Delta e_{jt} = f_j + g_t + \gamma \, dstart_{jt} + \delta \, dend_{jt} + \beta \, d_{jt} + \varepsilon_{jt}. \quad (1)$$

If firms commit to their disclosure policy independently from their private information and policy breaks are due to extraneous circumstances unrelated to earnings, an auto-regressive process would explain earnings in the next period ($\delta = 0$ and $\gamma = 0$). In contrast, we observe that a break in disclosure reduces by about a third how lagged earnings affect current earning expectations (see Table 7, Panel B). In other words, the act of ending (starting) a disclosure streak is associated with a significant decline (increase) in earnings.

Fact 3: Effect of non-disclosure spells

Considered jointly, the previous two facts suggest that the dependence of disclosure on earnings is affected by the history of disclosure. For instance, a firm setting a precedent for disclosing in the past may find it more difficult to withhold a forecast in the future and, as a consequence, its disclosure policy may become more sensitivity to earnings information.

To capture this effect, we regress current disclosure on earnings and the interaction of earnings with lagged disclosure:

$$d_{jt} = f_j + g_t + \beta_0\, e_{jt} + \beta_1\, d_{t-1} + \gamma\, e_{jt} d_{t-1} + \varepsilon_{jt}. \tag{2}$$

Consistent with the notion that reputation effects may affect a firm disclosure behavior we find that the disclosure-earnings sensitivity is significantly higher when the firm disclosed in the previous period compared to when it did not, or $\gamma > 0$ (see Table 7, Panel C). This analysis suggests that disclosure behavior is more sensitive to earnings when the firm has a recurrent pattern of disclosure. This is consistent with disclosures being more sensitive to earnings when the firm has established a reputation to disclose frequently and reinforces the need to incorporate dynamic effects into our structural model.

3 Theoretical model

This section extends the model in Dye (1985) and Jung and Kwon (1988) to a dynamic setting in which the manager has forward-looking motives and public information flows interact with management disclosure choices. We first lay out the theoretical framework and then discuss additional adjustments to the theory that must be made to accommodate some of the key stylized facts.

Each firm is traded by a large number of risk-neutral investors, referred to as the market. Time is discrete and indexed by $t = \{1, 2, ..., \infty\}$. Following Benmelech, Kandel and Veronesi (2010) and Beyer and Dye (2012), managers maximize the discounted value of a firm's stock price. That is, the manager's expected utility in period t is

$$\mathcal{U}_t = \mathbb{E}_t(\sum_{k=t}^{\infty} \beta^{k-t} P_k), \tag{3}$$

where $\beta \in (0, 1)$ is a subjective discount factor which we shall interpret as the rate at which managers sell shares or are exposed to share prices via vesting schedules in compensation arrangements (Edmans, Fang and Lewellen 2013; Marinovic and Varas 2017), P_k is the market price of the firm, and $\mathbb{E}_t(.)$ is the expectation at the beginning of period t.

The timeline in each period has three event dates, indicated by $t.1$, $t.2$ and $t.3$ in Figure 2. At the start of the period $t.1$, the market observes an early public signal $\hat{c}_t$ about end-of-period earnings e_t. Empirically, we will set $\hat{c}_t$ as the consensus analyst forecast but we can think of $\hat{c}_t$ more generally as a sufficient statistic capturing all new signals available to investors prior to a firm's disclosure.

At date $t.2$, the manager may receive information in the form of a private signal $\hat{s}_t$, which we denote as an indicator variable $\theta_t = 1$. Conversely, if $\theta_t = 0$, the manager does not receive material information beyond what is already contained in $\hat{c}_t$ and in prior public information. The market does not know the realization of θ_t and forms expectations given by a probability $p_t = 1 - \mathbb{E}_t(\theta_t)$. If information is received, the manager may decide to voluntarily disclose $\hat{s}_t$, that is, issue a forecast $d_t = \hat{s}_t$ about end-of-period earnings. By convention, $d_t = \emptyset$ indicates that no forecast is made. Then, the price P_t forms, reflecting the value of future earnings discounted at an objective market rate of return r and conditional on all public information $\mathcal{H}_{t-1}$ up to the end of period $t-1$ as well as any new information $\{\hat{c}_t, d_t\}$,

$$P_t = \mathbb{E}(\sum_{k=t}^{\infty} \frac{e_k}{(1+r)^{k-t}} | \mathcal{H}_{t-1}, \hat{c}_t, d_t). \tag{4}$$

At date $t.3$, the reporting period ends and the firm releases its earnings e_t. These earnings are publicly observed by investors and the public information set is updated to $\mathcal{H}_t = \{\mathcal{H}_{t-1}, \hat{c}_t, d_t, e_t\}$.

Following Dye (1985), the manager's information endowment $\theta_t \in \{0,1\}$ is random. We use here annual forecasts in our empirical setting given that Dye emphasizes such forecasts as a primary motivation for the model, e.g., "it is commonly believed that managers possess information about the firms they run, such as annual earnings' forecasts, whose release would affect the prices of their firms" (p.124). Over a horizon of up to a year, managers may be informed, or may not know more than the market (Chen, Goldstein and Jiang 2006) but whether they have received information is not known to outsiders or provable.

To port this model into a dynamic setting, we assume that there is inherent correlation in business transactions. We observe that earnings realizations exhibit a significant amount of serial correlation and assume that this creates correlation in the flow of information. Specifically, as in Einhorn and Ziv (2007), managers who have received information are more likely to face similar transactions in the future that will also yield advance information. In formal terms, the manager's information endowment θ_t is a hidden Markov chain with a transition matrix

$$\Pi = \begin{pmatrix} 1-\lambda_1 & \lambda_1 \\ \lambda_0 & 1-\lambda_0 \end{pmatrix}, \tag{5}$$

where $\lambda_0 \equiv \Pr(\theta_{t+1} = 0 | \theta_t = 1) \in (0,1)$ denotes the probability of moving from the informed to the uninformed state, and $\lambda_1 \equiv \Pr(\theta_{t+1} = 1|\theta_t = 0) \in (0,1)$ denotes the probability of moving from the uninformed to the informed state. The information endowment is persistent when becoming uninformed is less likely than remaining uninformed, or $\lambda_0 < 1 - \lambda_1$. We will sometimes parametrize the process θ_t by using the long-run probability of being uninformed $\overline{p}$ as well as the persistence of information endowments $\overline{m}$,

$$\overline{p} = \frac{\lambda_0}{\lambda_0 + \lambda_1}, \qquad \overline{m} = 1 - \frac{1}{2}(\lambda_0 + \lambda_1). \tag{6}$$

The process of earnings, consensus, and manager signal jointly satisfies (i) $e_t = \rho e_{t-1} + u_t$, (ii) $\hat{c}_t = e_t + v_t$ and (iii) $\hat{s}_t = e_t + w_t$, where $\varepsilon_t = (u_t, v_t, w_t)'$ is an iid normal vector with variance-covariance matrix $diag(\tau_u, \tau_v, \tau_w)^{-1}$, so that earnings e_t follow an AR(1) process and the information observed by investors and the manager $\hat{c}_t$ and e_t are noisy orthogonal signals about end-of-period earnings.[1] Note that the model predictions are entirely invariant to the mean of these processes, so we normalize all signals to zero mean. In the empirical implementation, we extract the residuals of the empirical processes from a fixed effects regression, so the processes used in the estimation have mean zero and match the model specification.

To fit empirical features of our setting, we make two adjustments to this framework. First, while it is convenient to state the model in terms of the vector of signals $(\hat{c}_t, \hat{s}_t)$, we observe empirically that analysts and managers communicate in terms of their expectation about future earnings. In the model, this means that we need to renormalize our variables to posterior expectations $c_t = \mathbb{E}(e_t|\hat{c}_t, \mathcal{H}_{t-1})$ as the model-implied market consensus and $s_t = \mathbb{E}(e_t|\hat{s}_t, \hat{c}_t, \mathcal{H}_{t-1})$ as the model-implied disclosure. From an informational perspective, there exists a one-to-one mapping between signals and posteriors, so that this transformation

[1] We assume that errors are uncorrelated. This is because, under sequential normal updating, correlation between errors is not separately identified from the noise in the signals. To see this, suppose the consensus forecast is $c = e + \varepsilon_c$ and the manager's forecast is $s = E(e|c, x)$, where the manager's private signal is $x = e + \varepsilon_m$. Suppose that all random variables are Gaussian and that the errors, ε_c and ε_m, are correlated. Then, by the projection theorem, we know that earnings can be represented as $e = s + \omega$ where ω is white noise. Hence, a model where the manager observes both the consensus and signal x is informationally equivalent to a model where the manager observes (the consensus and) a conditionally uncorrelated signal, s.

has no consequence on equilibrium payoffs and disclosure choices. When restating the model in terms of expectations, the stochastic process of earnings and signals becomes

$$\begin{pmatrix} e_t \\ c_t \\ s_t \end{pmatrix} = \rho \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} e_{t-1} + \begin{pmatrix} 1 \\ \frac{\tau_v}{\tau_u+\tau_v} \\ \frac{\tau_v+\tau_w}{\tau_u+\tau_v+\tau_w} \end{pmatrix} u_t + \begin{pmatrix} 0 \\ \frac{\tau_v}{\tau_u+\tau_v} \\ \frac{\tau_v}{\tau_u+\tau_v+\tau_w} \end{pmatrix} v_t + \begin{pmatrix} 0 \\ 0 \\ \frac{\tau_w}{\tau_u+\tau_v+\tau_w} \end{pmatrix} w_t. \tag{7}$$

Second, in our sample, firms that forecast frequently coexist with firms that have never made forecasts, suggesting that there is firm-specific heterogeneity about propensities to forecast. Unfortunately, with a limited time-series per firm, we cannot perfectly observe whether a firm never makes forecasts or did not make a forecast over the entire sample because it withheld them strategically. In other words, in steady-state, the market knows the entire history of forecasts and would identify which firms never disclose but the sample observed by the econometrician is contaminated by an extensive margin of firms who follow a non-forecasting model. We assume, therefore, that the strategic model applies to a fraction ξ of firms, which we estimate, while other firms are known to have a long history of not forecasting and never forecast. To estimate this model, we need to use simulated method of moments since a firm that did not issue a forecast in the sample could be a strategic firm choosing to conceal in all periods.

In a Markov perfect equilibrium (MPE), the payoff relevant public information is given by $z_t = (e_{t-1}, c_t)$ and p_t, and the information of the manager is (z_t, θ_t, s_t). From this point onwards, unless needed for clarity, we omit the time subscripts and refer to unrealized end-of-period variables using the $'$ notation.

For any public state (p, z), let $D(p, z) \equiv \{s \in \mathbb{R} | d(p, z, s) = 1\}$ be the manager's disclosure set when the manager disclosure strategy is an indicator variable $d(p, z, s)$. Let $P^D(z, s)$ and $P^{ND}(p, z)$ be the market price conditional upon disclosure and non-disclosure, respectively.

We require prices to be consistent with Bayes' rule and the manager's disclosure strategy:

$$P^D(z,s) = \frac{1+r}{1+r-\rho}\mathbb{E}(e'|z,s), \tag{8}$$

$$P^{ND}(p,z) = \frac{1+r}{1+r-\rho}\frac{pE(e'|z) + (1-p)E(e'\mathbf{1}_{D^c(p,z)}|z)}{p + (1-p)E(\mathbf{1}_{D^c(p,z)}|z)}. \tag{9}$$

In equation (9), as in Jung and Kwon (1988), $P^{ND}(p,z)$ is a weighted average between the payoff if the manager is uninformed $E(e'|z)$ and if the manager is informed but strategically withholding $E(e'\mathbf{1}_{D^c(p,z)}|z)$.

This price function assumes that prices are given by the expected present value of the firm's economic earnings, depending upon the AR(1) persistence parameter ρ and the manager subjective discount factor β. The market reassesses the probability that the manager was informed in the current period on the basis of earnings realization e'. Conditional on non-disclosure and earnings realization e', the updated probability that the manager will be uninformed in the next period is given by

$$p' = \varphi(p,z,e') \equiv \frac{p(1-\lambda_1) + (1-p)\lambda_0 E(\mathbf{1}_{D^c(p,z)}|e')}{p + (1-p)E(\mathbf{1}_{D^c(p,z)}|e')}. \tag{10}$$

When the manager withholds their signal, they retain an informational advantage about the firm's fundamentals and their expected information endowment. Specifically, investors do not know whether the manager strategically withheld unfavorable information or was uninformed. By contrast, if the manager discloses their signal, the market learns that the manager was informed and updates the probability that the manager will be uninformed in future to the probability λ_0.

We can now define the manager's optimization problem, when informed, as

$$V_1^D(p,z,s) = P^D(s,z) + \beta E\Big[(1-\lambda_0)V_1(\lambda_0, z', s') + \lambda_0 V_0(\lambda_0, z')\Big| z, s\Big] \tag{11}$$

$$V_1^{ND}(p,z,s) = P^{ND}(p,z) + \beta E\Big[(1-\lambda_0)V_1(p', z', s') + \lambda_0 V_0(p', z')\Big| z, s\Big] \tag{12}$$

$$V_1(p,z,s) = \max_{d\in\{0,1\}} \left[dV_1^D(p,z,s) + (1-d)V_1^{ND}(p,z,s)\right] \tag{13}$$

$$V_0(p,z) = P^{ND}(p,z) + \beta E\Big[\lambda_1 V_1(p', z', s') + (1-\lambda_1)V_0(p', z')\Big| z\Big], \tag{14}$$

where $V_1^{ND}(p,z,s)$ (resp.,$V_1^D(p,z,s)$) is the value function for an informed manager conditional on withholding (resp., disclosing), $V_1(p,z,s)$ is the informed manager's value function prior to making a disclosure choice, and $V_0(p,z)$ is the value function of an uninformed manager.

The market posterior expectation evaluated at alternative information sets is given by

$$e'|z \sim \mathcal{N}\left(c, \frac{1}{\tau_u+\tau_v}\right), \quad e'|z,s \sim \mathcal{N}\left(s, \frac{1}{\tau_u+\tau_v+\tau_w}\right), \tag{15}$$

$$s|z \sim \mathcal{N}\left(c, \frac{1}{\tau_u+\tau_v}\right), \quad s|c,e' \sim \mathcal{N}\left(e', \frac{\tau_w}{(\tau_u+\tau_v+\tau_w)^2}, \right). \tag{16}$$

We can now formally define the equilibrium concept.

Definition 1. *A Markov Perfect Equilibrium (MPE) is a tuple $\langle P^D, P^{ND}, d, \varphi, V_0, V_1, V_1^D, V_1^{ND}\rangle$, such that*

1. *The market price is*

$$P = \begin{cases} P^{ND}(p,z) & \text{if } d(p,z,s) = 0 \\ P^D(z,s) & \text{if } d(p,z,s) = 1, \end{cases} \tag{17}$$

where P^D and P^{ND} are given by (8) and (9).

2. *The disclosure strategy* $d(p,z,s) \in \{0,1\}$ *is*

$$d(p,z,s) \in \arg\max_{d\in\{0,1\}} \left[dV_1^D(p,z,s) + (1-d)V_1^{ND}(p,z,s) \right], \tag{18}$$

if the manager is informed and is $d(p,z,s) = 0$ *if the manager is uninformed.*

3. *The evolution of market beliefs is*

$$p' = \begin{cases} \varphi(p,z,e') & \text{if } d(p,z,s) = 0 \\ \lambda_0 & \text{if } d(p,z,s) = 1, \end{cases} \tag{19}$$

where $\varphi(p,z,e')$ *is given by* (10).

4. *The value function of the informed manager solves*

$$V_1(p,z,s) = \max\left\{V_1^D(z,s), V_1^{ND}(p,z,s)\right\}, \tag{20}$$

where

$$V_1^D(z,s) = P^D(s,z) + \beta E\Big[(1-\lambda_0)V_1(\lambda_0,z',s') + \lambda_0 V_0(\lambda_0,z')|z,s\Big] \tag{21}$$

$$V_1^{ND}(p,z,s) = P^{ND}(p,z) + \beta E\Big[(1-\lambda_0)V_1(p',z',s') + \lambda_0 V_0(p',z')|z,s\Big]. \tag{22}$$

5. *The value function of the uninformed manager solves*

$$V_0(p,z) = P^{ND}(p,z) + \beta E\Big[\lambda_1 V_1(p',z',s') + (1-\lambda_1)V_0(p',z')|z\Big]. \tag{23}$$

Below we develop the intuition for the economic tradeoffs that drive the equilibrium. Consider the manager's disclosure and withholding incentives. Withholding information carries two benefits from the manager's standpoint. First, by hiding bad news, the manager delays the decline in stock price because the market is uncertain about the true cause of a non-disclosure. Because the manager benefits from higher short-term stock prices, such a delay is attractive. Second, the manager further influences the market's perception about their future information endowment. By pretending to be uninformed, the manager can increase the perceived probability that they will be uninformed in the future. This, in turn, mitigates the price penalty triggered by non-disclosure and increases the option value from continuing to withholding information. Withholding information thus entails a reputational benefit.

Naturally, when the information endowment is iid, namely, $\lambda_0 = 1 - \lambda_1$, the reputation benefit of withholding is absent. In this case, market beliefs are constant and independent of the manager's disclosure choices. For this reason, the manager's disclosure strategy collapses to that of a static model. We provide this observation as a benchmark.

Proposition 2. *When the information endowment is iid, i.e., $\lambda_0 = 1 - \lambda_1$, there is a unique equilibrium where in each period the manager uses a myopic threshold τ, defined as*

$$f(\tau; \frac{\lambda_0}{\lambda_1}) = \int_{-\infty}^{\tau} \Phi(x)dx + \frac{\lambda_0}{\lambda_1}\tau = 0. \tag{24}$$

In each period, the manager discloses when observing $\frac{s_t - E[s_t|z_t]}{\sqrt{Var(s_t|z_t)}} \geq \tau$.

The optimal disclosure strategy in the static model can be obtained as a threshold $\tau < 0$ on the standardized signal s_t, such that the manager discloses any standardized signal above the threshold. Notice that the threshold depends only on the information endowment process. In the static case, the frequency of disclosure is independent of most firm characteristics, including the precision of the manager's signal τ_w, the precision of consensus signal τ_c, and the volatility of firm fundamentals σ_u. Hence, these characteristics of the information environment must affect disclosure behavior via dynamic channels.

To gain additional intuition on the effect on information endowment in the iid model, denote $\ell = \lambda_0/\lambda_1$. Then, it follows from the implicit function theorem and equation (24) that

$$\frac{\partial \tau}{\partial \ell} = -\frac{f_\tau}{f_\ell} = -\frac{\tau}{\Phi(\tau) + \ell} > 0,$$

implying that the strategic withholding becomes more likely when the probability that the manager is uninformed increases (Dye 1985; Jung and Kwon 1988).

We consider next the structure of the manager's payoffs after a disclosure, obtaining a simple formula linking manager value in the case of disclosure to the implied market price.

Proposition 3. *In any equilibrium,* $V_1^D(p, z, s) = \frac{P^D(z,s)}{1-\rho\beta}$.

The manager's payoff conditional on disclosure is linear in both public information z and the manager's private signal s. On the surface, this property seems to ignore the option value of withholding information in future periods: the disclosing manager's payoff is the same as if the manager had committed to full disclosure forever. Upon disclosure, the manager's and market's information sets coincide so that there are no further channel for the manager to affect expected prices.

The manager's payoff given non-disclosure also increases in the value of their signal s because a higher s has a positive expected reputation effect. Since higher values of s are correlated with higher realizations of earnings e, the market is more likely to believe the manager was uninformed conditional on non-disclosure. However, the payoff given non-disclosure is non-linear in s. The existence of a threshold equilibrium is therefore not guaranteed. Indeed, when the information endowment is persistent and the manager's signal is very precise, there is no threshold equilibrium in our game. We show this result in a special case of the model with no public signal z and iid earnings.

Proposition 4. *Assume (i) the information endowment is persistent,* $\lambda_0 < 1 - \lambda_1$*, (ii) there is no public signal* z *(* $Var(n_t) \to \infty$*), (iii) the manager is almost perfectly informed (*$Var(v_t) \to 0$*), and (iv) the earnings process* e_t *is iid. Then, there is no equilibrium, such that for any current market belief* p*, the manager adopts a threshold equilibrium, defined as disclosing if* $s > k_p$ *and withholding if* $s < k_p$*.*

A standard threshold structure characterizing the equilibrium in the static game is not assured in our dynamic setting. The reason behind Proposition 4 follows Grubb (2011), who shows that equilibrium disclosure strategies are mixed in a dynamic model with reputation concerns. In our setting, we can understand this property by assuming that the firm does use a threshold equilibrium and considering two informed managers, A and B, whose signals lie slightly below and slightly above a conjectured threshold, respectively. Manager A must prefer to withhold; hence, the net effect of withholding on the current price and on the reputation is non-negative. Now, suppose that manager B deviates to withhold information, as well. This will cause almost the same current price effect; however, the reputational effect is different. After earnings are revealed, the market will attribute the above-threshold earnings to an uninformed manager with probability one because this is the updated belief on the equilibrium path. Hence, manager B benefits more from withholding more than manager A does.

A key to this argument is the fact that reputations are endogenous and depend on market expectations. The reputational effects are stronger near a disclosure threshold because this is where the market learns the most from realized earnings, and the same effects are also stronger for managers with more favorable information because the market makes more favorable reputational updates after positive earnings are reported. Fortunately, some amount of noise in realized earnings can preserve threshold equilibria. To address this potential issue, we solve the model without assuming threshold equilibrium and numerically assess whether each type is choosing to disclose optimally. For levels of noise in earnings consistent with our empirical sample, we find that the optimal policy is a threshold equilibrium.

The comparative statics of the model cannot be easily recovered analytically but we can plot below the comparative statics in the steady state of the model, with parameters $\sigma_u = \sigma_v = \sigma_v = 1$ and $\beta = 1/(1 + r) = .9$. Figure 3 reports the probability of withholding as a function market beliefs for several levels of persistence in the manager's information endowment. The model with iid information endowment coincides with the static or myopic solution (dotted curve), and the other lines represent equilibria in which the persistence is set to higher levels. As the information endowment becomes more persistent, the reputational benefit of withholding increases and, therefore, the manager becomes more likely to withhold for any given market belief. For all levels of persistence, managers are more likely to withhold information strategically when the market belief p is high because the market penalty for non-disclosure becomes less severe.

We elaborate next on the mechanisms at play in the dynamic model by considering how the economic primitives affect the solution of the model. Managers' disclosure choices are affected by several aspects of the information environment. First, disclosure choices depend on the extent to which the manager cares about the current price. A higher subjective discount rate β is equivalent to an increase in the reputational concern. Figure 4 plots the level of strategic withholding in the model as a function of market beliefs for several values of impatience. A more patient manager tends to withhold information more often to affect market perceptions about θ, thus foregoing a short-term price benefit to gain disclosure flexibility in the future. In the limit, an infinitely patient manager would wait for the realization of all uncertainties via public news and earnings realizations.

The distribution of the manager information endowment θ is also a key determinant of disclosure behavior. As noted above, it is convenient to characterize the distribution of θ with two parameters: the long-run probability of being informed $\bar{p} \equiv \frac{\lambda_0}{\lambda_0+\lambda_1}$ and the persistence of information endowment $\overline{m} \equiv 1 - \frac{\lambda_0+\lambda_1}{2}$. The effect of changes in $\bar{p}$ is intuitive: the higher the probability that the manager is informed, the more skeptical is the market about the firm value when the manager does not disclose, which, in turn, prompts more disclosure, as in the static case.

The effect of persistence $\overline{m}$ is the combination of two effects. Recall from Figure 3 that, with more persistence, a manager discloses less for any given belief about the likelihood of being uninformed. But persistence also increases learning over time, changing the underlying steady-state distribution of market beliefs p in the long run as well. We illustrate this effect in the top panel of Figure 5, which shows that the dispersion in posterior beliefs increases with greater persistence. Investors know more about the information endowment when it is more persistent, thus reducing the amount of equilibrium uncertainty about information endowment. In turn, with less uncertainty, the steady-state disclosure strategies tend to be closer to unravelling. In the limit case of a fully persistent information endowment, markets will know the information endowment almost perfectly in the long run, implying that all informed managers, because they are known to be informed by the market, will disclose with probability one. Put differently, with sufficient persistence, reputations are not renewed and will dissipate in the long run (see also Cripps, Mailath and Samuelson 2004).

When the market learns about the information endowment quickly, there is less value to be gained from strategic concealment and, consistent with unraveling theory (Grossman and Hart 1980; Milgrom 1981), an informed manager discloses more often. We plot in Figure 5, bottom panel, the steady-state probability of disclosure which results from the joint effect of persistence on strategic concealment at any given belief and quicker learning. The greatest level of concealment in steady-state is achieved for intermediate levels of uncertainty as these preserve the highest level of uncertainty in the dynamic game.

We also consider the effect of the variance in fundamentals σ_u. First, observe that if earnings were not observable, as in the static version of the model, the volatility of fundamentals would not affect the probability of disclosure: this variance would simply scale the manager's payoffs without affecting the relative trade-offs. However, observable earnings play a disciplining role in the manager's disclosure behavior because they allow investors to refine their beliefs regarding whether the manager has concealed information in the past. Recall that the manager bears a current price decrease when concealing information, in anticipation of building a reputation. However, in environments with less noise in earnings, realized earnings reveal the private information known to the manager, and thus the reputational benefit is weakened. As illustrated in Figure 6, a manager withholds less for a given market belief when the variance of fundamentals is low.

4 Data and Estimation

Our sample is from the I/B/E/S earnings announcement database for fiscal years ending between January 1st, 2004 and December 31st, 2016 and firms traded in a major US exchange. The details of the sample selection are reported in Table 2. We start the sample in 2004 because of significant changes in US regulation in 2000 and 2002, which have altered both the incentives to disclose information and the collection of forecasts. Since August 2000, Regulation Fair Disclosure (Reg FD) prohibits most private communications between managers and market analysts. By shutting down this communication channel, Reg FD appears to have increased the frequency of public managerial forecasts. Also, since July 2002, the Sarbanes-Oxley Act (SOX) dramatically increased internal controls and management responsibilities. From a data-collection perspective, SOX also required conference calls to be in transcript form, allowing for convenient identification of manager forecast disclosure.

We construct a sample of raw earnings per share (EPS). Earnings in I/B/E/S are reported as pro-forma earnings calculated under the same accounting principles as analysts' and management forecasts. The initial sample includes $83,935$ firm-years from $10,542$ unique firms. We match to Compustat and CRSP, non-missing and non-negative equity, and require non-missing announcement and lagged announcement dates to create a window for management forecasts.

We obtain management forecasts from the I/B/E/S management forecast guidance (CIG) database. Managers may make forecasts on a variety of items (earnings, revenue, gross margins, etc.) and for different time horizons, in press releases or conference calls. The most common forecasts are one-quarter earnings forecasts; however, these tend to occur at a point where most of the uncertainty has already realized and are typically on a schedule. Hence, we focus on one-year ahead earnings forecast, which are the next most common type of forecast. The majority of forecasts are made bundled with the prior earnings announcement, that is, are made in a single press release including both current earnings and the forecast for next-year earnings. In our final sample, such bundled forecasts form 90.2% of the forecasts, which implies a fairly consistent forecasting horizon of between 10 and 11 months before a fiscal year end.

We merge them with the I/B/E/S earnings database by using I/B/E/S unique tickers and the forecast period end date, retaining only annual forecasts where the forecast period end date can be matched to the I/B/E/S earnings announcement sample. This selection yields a sample of 66,602 forecasts. We require a forecast to be made after the prior earnings announcement date but at least six months before the period end date. We remove forecasts made after the fiscal year end because they are less likely to be consistent with a model of incomplete information endowment. For periods with multiple forecasts, we use the earliest forecast, which yields a sample of 13,513 forecasts. In the final sample, 90.2% of the forecasts were made during the earnings announcement.

We calculate the raw forecast, as it was made, by multiplying the forecast in I/B/E/S, adjusted for the number of shares with the I/B/E/S adjustment factor. For each firm-year with or without a forecast, we require a measure of market expectation about realized earnings. For any firm-year with a forecast, we use the I/B/E/S CIG analyst consensus; this number is provided adjusted for stock splits by I/B/E/S, and we unadjust it by multiplying by the adjustment factor. This consensus measure reflects the consensus before a management forecast is made. We calculate a consensus by using the I/B/E/S analyst file, which reports all annual forecasts made by financial analysts. The consensus is defined as the median of the last five unique analyst forecasts made prior to the forecast or, in the absence of a forecast, prior to the earnings announcement. This implies that the consensus is always constructed from analyst forecasts made prior to the management forecast. We remove observations for which we are unable to form a market expectation and drop observations with missing assets or earnings, implying a sample of 54,167 unique earnings announcements in 8,109 unique firms; during this period, a total of 13,309 management forecasts were made. Lastly, we winsorize all variables at the 1% level and scale realized earnings, forecasts and consensus by assets, in order to control for changes in the size of firms that are not within the scope of our analysis.

Table 3 reports several descriptive statistics. The overall disclosure frequency in the sample is about 23.5%, suggesting that managers commonly provide numerical forward-looking information. About half of the forecasts lie between 4% and 10% of assets. Consistent with the theory and the stylized facts, forecasts tend to be greater than actual earnings, with a median forecast of 6.21% of assets versus median earnings of 3.90%. Forecast errors in our sample are small, and the average bias was .3% of assets in our sample, thus suggesting that the difference in earnings and forecasts is not driven by biased reporting. In particular, this fact suggests that selection, rather than reporting bias, is the main factor driving observed forecasts. It is also worth noting that most firms in our sample are medium or large-sized, with median total assets of $1.4 billion and median market capitalization of $1 billion.

We fix the values of the parameters of the quantitative dynamic model in two main steps. First, we calibrate five of the parameters (β, r, ρ, σ_u, σ_v) externally, drawing on both comparable quantitative exercises from the literature on dynamic corporate finance as well as a moment-matching exercise for the parameters of some purely exogenous series. Second, conditional upon these calibrated parameters, we estimate each of the three disclosure-related parameters (σ_w, λ_0, λ_1). Because the theoretical model rules out non-strategic behavior by assumption for any firm with disclosure breaks, we estimate an additional parameter ξ capturing the probability that a manager is non-strategic and always discloses when informed, i.e., $d_t = \theta_t$. In addition, to fit the large number of firms that never disclose, we estimate a probability ζ that a firm never discloses. In summary, a firm in our model may strategically withhold with probability $1 - \xi - \zeta$, always discloses when informed with probability ξ and never discloses with probability ζ. The estimation is a straightforward application of simulated method of moments, in which we simulate a dataset in which a firm is drawn to be one of the three possible types.[2]

[2] We also estimated the model with either ξ or ζ were zero, but found that the two types of heterogeneity were required to explain disclosure behavior. Intuitively, it is statistically unlikely that a firm that never discloses or always discloses is driven by the same frictions as firms with time-varying disclosures and, indeed, estimating a model without this heterogeneity tends to predict implausibly high levels of persistence.

We assume a model period of one year, corresponding to the frequency of our I/B/E/S data. We set the value of the objective discount rate to $r = 4\%$, approximately in the middle of a range of annual discount factors commonly used in dynamic corporate finance.[3] We set the manager's subjective discount rate to $\beta = 1 - (1/3.29)$ based on the median vesting duration in Gopalan, Milbourn, Song and Thakor (2014). This choice is consistent with myopia induced by vesting practices in Edmans, Fang and Lewellen (2013); in robustness checks, we also report estimates of the model when calibrating impatience to the median CEO tenure in Taylor (2010) of 6 years.

The values ρ, σ_u, and σ_v govern the persistence and volatility of the fundamental earnings process, as well as the signal precision for analyst forecasts. The realized earnings process e_t and consensus forecasts c_t are directly observable in our data for all years. Both series are exogenous to the model disclosure choice, so to economize on the number of computationally costly estimated parameters, we fix the values of these parameters by matching three data moments directly. Table 4 reports the values of the three calibrated parameters and the targeted moments in the model and the data. The model closely reproduces these targeted moments, and the resulting parameter values reveal substantial persistence in fundamental earnings as well as a large amount of noise in the analyst signal relative to underlying earnings. Although the mapping between moments and parameters here is joint and not one-to-one, the autocorrelation of earnings is informative for the value of persistence ρ, and the dispersion of earnings and analyst forecast errors help pin down the values of σ_u and σ_v, respectively.

Armed with values for each of the other parameters in the model, we now turn to the estimation of the structural parameters $(\xi, \zeta, \sigma_w, \lambda_0, \lambda_1)$. These five parameters relate directly to the unobserved series of manager signals and information endowments in the model, and we cannot observe either of these series directly in the dataset. Instead, we turn to a SMM procedure. We minimize the sum of squared deviations between the model and data values

[3] See, for comparison, the dynamic firm-level investment model calibrations studied in Cooper and Ejarque (2003), Hennessy and Whited (2005) and Hennessy and Whited (2007).

of a selected set of moments relating to the disclosure process.

To ensure identification, moment-based structural estimation strategies rely crucially on the selection of an appropriate set of targeted moments that are informative for the underlying parameters of interest. To guide our intuition about which moments to draw on for identification, we consider how the model would be estimated within a myopic model. In this model, the disclosure rate and dynamics rely directly on parameters λ_0 and λ_1 of the manager information endowment. This dependence reflects a direct effect through the information endowment itself, as well as an indirect effect through the dynamics of the market belief p surrounding the manager's information state at any point in time. To capture this relationship, we target the frequency of disclosure, as well as the probability of disclosure given prior disclosure in each of the past three years over firms with time-varying disclosure. In addition, note that conditional upon disclosure, the dispersion of manager forecast errors depends directly on the precision of manager signal σ_w. Therefore, we target the standard deviation of the manager forecast errors in the data. Lastly, to identify ξ and ζ, we incorporate in our estimation the fraction of firms that always disclose and the fraction of firms that never disclose.

With the targeted moments from the data in Table 5 and given any candidate values for parameters σ_w, λ_0, and λ_1, we compute comparable model moments directly from the stationary distribution of the model after solving the model numerically. More information on the solution technique we use for the model is available in the appendix. We numerically minimize the SMM objective function, the sum of squared deviations between data and model moments, to obtain point estimates. The standard errors of the estimated parameters follow standard SMM formulas and rely upon a bootstrap procedure to calculate the empirical moment covariance matrix.[4]

[4]In particular, we minimize the SMM objective by using particle swarm optimization, a numerical global stochastic optimization routine. Given point estimates, we compute the covariance matrix of the targeted moments from the data Ω via a block bootstrap procedure, resampling at the firm level and accounting for within-firm serial correlation. We then compute an approximate value of the Jacobian M of model moments with respect to the estimated parameters by using forward numerical differentiation from the point estimates and averaging over step sizes of .25%, .5%, and 1%. With these objects in hand, the asymptotic

The final column of Table 5 reveals that the estimated model reproduces the targeted moments well. We cannot expect an exact match between data and model moments given our nonlinear model and the overidentified nature of the estimation exercise with three parameters and five targeted moments. However, the estimated model closely reproduces the unconditional disclosure probability of around 60%, the high persistence of disclosure observed in the data, and the dispersion of manager forecast errors. We also conduct a J-test of over-identifying restriction, which rejects that all moments are jointly met at conventional significance level. While this is not surprising given that we use a stylized model to fit the entire cross-section, this seems to be in part driven by $\mathrm{P}(d_t = 1 \mid d_t \text{ break})$ which is 43% in the data but 54% in the model. This fact suggests that there are likely additional institutional forces driving disclosure choices. For example, there could be a more complex pattern of unobserved heterogeneity in manager characteristics, information endowments or other drivers for withholding (e.g., proprietary costs, differences in internal controls or investor demands). This is a natural caveat to our analyses but also suggests additional forces to fully explain the data moments.

Table 10 reports the parameter estimates. We estimate substantial persistence in manager information endowments with only around 4% probability of future information loss given a manager signal today ($\hat{\lambda}_0$) and less than 6% chance of becoming informed for a currently uninformed manager ($\hat{\lambda}_1$). The iid hypothesis of static information endowments $\lambda_0 = 1 - \lambda_1$ is rejected at the 1% significance level. Indeed, our point estimate of persistence is $\overline{m} = 1 - \frac{1}{2}(\lambda_0 + \lambda_1) = .95$. The high degree of persistence confirms the importance of forward-looking considerations we emphasized above. The estimated probability of being uninformed, conditional on being a strategic firm, is $\overline{p} = \lambda_0/(\lambda_1 + \lambda_0) = 40\%$, implying that managers possess one-year ahead private information that they might credibly disclose 60% of the time. Also, the manager signal dispersion $\hat{\sigma}_w = .0154$ implies more precise information for managers than analysts, intuitive in this context given their location within the firm.

distribution of the estimates $\hat{\theta}$ of the underlying parameters θ is given by $\sqrt{N}(\hat{\theta} - \theta) \rightarrow N(0, \Sigma)$, where $\Sigma = (M'M)^{-1} M'\Omega M (M'M)^{-1}$ and N represents the number of firms.

5 Model validation

In this section, we revisit the three stylized facts discussed in Section 2, within the quantitative framework of the model and ask whether the model appears to capture the salient stylized facts. These moments are untargeted, that is, they are not directly matched in the estimation procedure. Untargeted moments provide additional confidence that the assumptions of the model capture key economic mechanisms, but they also point out to aspects of the data that are not well captured in the model.

In Figure 7, we plot the empirical earnings as a function of the disclosure date versus the model implied earnings. The model fits reasonably well the increase in earnings at the disclosure date and, from serial correlation, earnings after the disclosure date, implying that the framework captures well strategic selection in forecasts, and what follows after the forecasting year (see Table 7, Panel A). Although the model is consistent with an increase in earnings prior to the forecast, this increase is greater in the data than it is in the model, so that the model under-predicts the increase in earnings prior to forecast. To understand why, we examine in the bottom panel, the consensus forecast around forecasting dates. According to the model (blue), the consensus should adjust primarily after a forecast is made, but we do find that analyst reflect earnings news even before the forecast is made. This suggests that there may be private information leakage about the forecast that may occur before the forecast is released (e.g., a manager may indicate that a forecast is forthcoming).[5]

Another result of our model is that disclosure breaks predict earnings. Specifically, we find that a firm that breaks its disclosure policy, e.g., from disclosure to non-disclosure, tends to realize lower earnings and vice-versa. We ask next whether the model captures well this important feature of the data and run the regression in (1) with data simulated from the model. The model reproduces the main directional effects obtained in the data, and all the coefficients are significant with the same sign as in the data.

[5]Recall that we compute the consensus using the period before a forecast is made, so that our consensus variable is not a function of the presence of a forecast unless analysts have private information in advance of the forecast.

Next, we revisit the ordinary least squares regression in (2) of disclosure on the interaction of short disclosure spells and earnings which captures the presence of reputation effects. The coefficient on the interaction of short spells (whether the firm disclosed in the prior period) with earnings is very similar between the model and the data. One difference is that, in the model disclosure lags do not explain much of the disclosure behavior (the coefficient is significant but low) once controlling for realized earnings; by contrast, the data shows that lagged disclosure remains a key variable. This may reveal that there is additional stickiness in the model that is not fully explained by the strategic incentives.

Note that our baseline model unambiguously rejects the hypothesis that the managers are non-strategic, since the estimated coefficient $1-\hat{\xi}-\hat{\zeta}$ is different from zero at any conventional level of significance. To explore the possibility of no strategic withholding further, we re-estimate the entire model with the restriction that no manager strategically withholds, i.e., $1 - \xi - \zeta = 0$. The targeted moments and estimated parameters of the restricted model are reported in column 4 of Tables 5 and 10. As expected from our baseline estimates, the restricted model fits the targeted moments significantly worse than the baseline, providing additional evidence consistent with strategic reporting. In addition, the non-strategic model does not explain well the untargeted empirical facts. The earnings and consensus dynamics in Figures 7 are flat and the stylized facts documented in Table 7, Panels B and C, are insignificant or with incorrect sign. Together, the targeted and untargeted moments provide strong evidence of the importance of strategic to explain observed disclosures.

6 Counterfactuals

Over the last two decades, policy makers in the US have encouraged firms to disclose forward-looking information. Indeed, in 1995, US regulators enacted the *Private Securities Litigation Reform Act* which includes a safe harbor provision that protects managers from litigation arising from unattained projections of forward-looking information. Following a similar logic, US regulators enacted in 2002 *Regulation Fair Disclosure* to encourage public disclosures and prohibit managers from selectively disclosing information to certain market participants while excluding others. Regulators' preference for more disclosure emphasizes the importance of quantifying the magnitude of firms' strategic withholding of information.

To measure the economic consequences of managerial forecasts, we compare our baseline estimates to three primary counter-factuals in the form of simple comparative statics exercises:

(i) no strategic concealment, holding all coefficients fixed but setting the disclosure threshold to $-\infty$ so that a manager always discloses when informed;

(ii) no informative forecasting, setting the noise in the manager's signal to ∞ while holding all other coefficients fixed;

(iii) no consensus forecast, setting the noise in the consensus to ∞, while all holding all other coefficients fixed.

A key moment of interest to evaluate the information flow to investors is the market pricing error after the forecast data, which we measure in Table 8 as the square root of the mean squared error (RMSE) between the market estimate of fundamental value by the time of the forecast and the true earnings. Within the context of normally-distributed random processes, it is an estimate of the standard-error of posterior expectations of earnings and capture how much the forecasting process reveals information to investors. In our baseline model, the RMSE is 6.3% of assets; without strategic withholding (i), the RMSE would be 6%. In relative terms, the RMSE decreases by 4% of the standard deviation of earnings innovations (7.5%).

This modest effect of strategic withholding is caused, in the model, by the disciplining effect of investor skepticism and is somewhat surprising given the importance of management forecasts. In a counterfactual with no managerial forecasts, we find that the relative RMSE would be 3% higher relative to the baseline. Interestingly, analyst forecasts contribute to discipline forecasting behavior and themselves play an important role as an information channel. Without a consensus, the relative RMSE would be 4% higher than in the baseline.

Note that these estimates do not imply that investors are made worse-off by withholding and, as such, the RMSE provide a measure of public information but not a measure of investor welfare. In fact, investors in the model trade the firm at its expected price conditional on all public information. In a broader context, however, undisclosed private information may disseminate to investors through other means and reveal itself from trades after public announcements (Manela 2014). Private information may also be valuable to informed agents adjusting their actions before uncertainty is publicly realized (Kadan and Manela 2018).

An additional feature of interest is the probability that an informed firm chooses to conceal information. In the benchmark estimates, this occurs with probability 26.1%. We also verify that managers with very noisy information are more likely to conceal, with a probability of strategic withholding of about 48.5%. We also observe that analyst consensus not only brings new information into the market but also makes managers less likely to conceal, reducing the probability of concealment to 21.1%.

7 Robustness Analyses

We examine, below, the robustness of our baseline analysis to changes in the parameters of the model, selecting for further inspection parameters that have pre-estimated from the earnings and analyst consensus data in Table 4 $(\rho, \sigma_u, \sigma_w)$, as well as the impatience of the manager β. Table 9 tracks the relative RMSE and the probability of concealment $P(ND|Informed)$ under strategic reporting and in the model with no concealment, against the benchmark values $(\rho, \sigma_u, \sigma_w, \beta) = (.306, .075, .138, .696)$.

Varying the parameters from 5% to 20%, we obtain estimates in the same range in sensitivity analyses, with the difference in relative RMSE between our baseline model and no-concealment varying between 7% and 10% versus 8% for the baseline estimates, and a probability of concealment varying from .24 to .29 versus .26 in the baseline. As expected, the serial correlation in earnings ρ has little effect on disclosure behavior because markets respond to the unexpected components of the earnings process. Increases in the noise in earnings σ_u and the noise in consensus σ_v increase the RMSE, because they reduce other sources of information. Varying β from $5/6 = .833$, corresponding to a 6-year median tenure in Taylor (2010), to .663, small changes to the discount factor of the manager do not seem to have large effects on the results.

While the main results are intended to describe effects summarizing the entire cross-section of firm disclosure, we consider next subsamples likely to exhibit heterogeneity in the drivers of corporate disclosure. Specifically, we classify firms as above or below the mean on three observable characteristics: number of analysts, size (market capitalization) and average analyst errors for estimates made before the management forecast. We choose these variables because they are associated to the occurrence of forecasts in the cross-section (Beyer, Cohen, Lys and Walther 2010) and may proxy for firm characteristics that affect the information environment. Naturally, our model does not allow us to draw causal implications about these variables because analyst following, capitalization or average errors could may be affected by the disclosure behavior and, therefore, these analyses are solely intended to assess additional variation that is not captured by the model.

The effect of concealment becomes more muted in the subsamples with a high number of analysts, high size or low analyst consensus errors, which intuitively suggests that these subsamples with competing information channels may feature other factors that limit the scope for strategic concealment. For these firms, the market consensus may already reveal a large portion of the information known to management, implying a small effect of concealment on RMSE. We find the largest changes in RMSE and probability of concealment in the subsamples with low number of analysts and low sizes, which likely indicate settings in which the management forecast is a primary source of information. The effect of concealment on RMSE is the largest in the subsample with high analyst consensus errors, consistent with analysts' information imperfectly substituting for the absence of a management forecast.

8 Conclusion

This paper develops a model of voluntary disclosure in which managers have forward-looking motives and their disclosure choices interact with both analysts' consensus forecasts and mandatory earnings announcements. Building on Dye (1985) and Jung and Kwon (1988), we assume that markets are uncertain about managers' information endowment. We show that small levels of persistence in the information endowment can fuel reputation concerns, capable of generating significant disclosure stickiness. This result accords with the notion originally advanced by the survey of Graham, Harvey and Rajgopal (2005) that managers' reluctance to disclose is driven by the fear of setting a disclosure precedent that might expose the stock price to negative market reactions during non-disclosure periods.

The persistence of information endowments has a subtle effect on disclosure policy. On one hand, persistence boosts managers' reputation concerns: by manipulating disclosure choices managers may attempt to influence market perceptions regarding their information endowment. By withholding information and thus pretending to be uninformed, a manager can increase the option value of withholding information in the future. The presence of this reputation effect makes managers more reluctant to disclose information. Persistence also has another effect, running in the opposite direction. Higher persistence leads, in the long run, to greater dispersion in the distribution of market beliefs about the manager's information endowment. On average, this belief dispersion causes a stronger unraveling pressure that induces extra disclosure. In some cases, the belief dispersion effect dominates the reputation effect.

Overall, the estimate suggests that voluntary channels filter out a measurable portion of the information available to managers, in contrast to the predictions of unravelling theory (Grossman 1981; Milgrom 1981). Yet, having focused on an environment in which information endowment is exogenous (Dye 1985; Jung and Kwon 1988), we know little of the potential mechanisms affecting disclosure frictions, and whether corporate choices ranging from investment to capital structure could affect the voluntary disclosure process. Building strategic models of the interaction between factors affecting production of information and their dissemination suggests rich avenues for future research.

References

Acharya, Viral V, Peter Demarzo, and Ilan Kremer, "Endogenous Information Flows and the Clustering of Announcements," *The American Economic Review*, 2011, *101* (7), 2955–2979.

Ball, Ray and Philip Brown, "An empirical evaluation of accounting income numbers," *Journal of accounting research*, 1968, pp. 159–178.

Benmelech, Efraim, Eugene Kandel, and Pietro Veronesi, "Stock-Based Compensation and CEO (Dis)Incentives," *The Quarterly Journal of Economics*, 2010, *125* (4), 1769–1820.

Bertomeu, Jeremy, Anne Beyer, and Daniel J Taylor, "From casual to causal inference in Accounting Research: The need for theoretical foundations," *Foundations and Trends® in Accounting*, 2016, *10* (2-4), 262–313.

_, Paul Ma, and Iván Marinovic, "How often do managers withhold information?," *Available at SSRN*, 2016.

Beyer, Anne and Ronald A. Dye, "Reputation management and the disclosure of earnings forecasts," *Review of Accounting Studies*, January 2012, *17* (4), 877–912.

_, Daniel A. Cohen, Thomas Z. Lys, and Beverly R. Walther, "The financial reporting environment: Review of the recent literature," *Journal of Accounting and Economics*, 2010, *50* (2-3), 296–343.

_, Ilan Guttman, and Iván Marinovic, "Earnings management and earnings quality: Theory and evidence," *Available at SSRN 2516538*, 2014.

Chen, Qi, Itay Goldstein, and Wei Jiang, "Price informativeness and investment sensitivity to stock price," *The Review of Financial Studies*, 2006, *20* (3), 619–650.

Cheynel, Edwige and Michelle Liu, "A Simple Structural Estimator of Disclosure Costs," 2017.

Cooper, Russell and Joao Ejarque, "Financial frictions and investment: requiem in q," *Review of Economic Dynamics*, 2003, *6* (4), 710–728.

Cripps, Martin W, George J Mailath, and Larry Samuelson, "Imperfect monitoring and impermanent reputations," *Econometrica*, 2004, *72* (2), 407–432.

Dye, Ronald A., "Disclosure of Nonproprietary Information," *Journal of Accounting Research*, 1985, *23* (1), 123–145.

Edmans, Alex, Vivian W Fang, and Katharina A Lewellen, "Equity vesting and managerial myopia," Technical Report, National Bureau of Economic Research 2013.

Einhorn, Eti and Amir Ziv, "Unbalanced Information and the Interaction between Information Acquisition, Operating Activities, and Voluntary Disclosure," *The Accounting Review*, 2007, *82* (5), 1171–1194.

— and — , "Intertemporal Dynamics of Corporate Voluntary Disclosures," *Journal of Accounting Research*, 2008, *46* (3), 567–589.

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

Frankel, Alex and Navin Kartik, "Muddled Information," Technical Report, Unpublished 2014.

Gopalan, Radhakrishnan, Todd Milbourn, Fenghua Song, and Anjan V Thakor, "Duration of executive compensation," *The Journal of Finance*, 2014, *69* (6), 2777–2817.

Graham, John R, Campbell R Harvey, and Shiva Rajgopal, "The economic implications of corporate financial reporting," *Journal of accounting and economics*, 2005, *40* (1), 3–73.

Grossman, Sanford J., "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law and Economics*, 1981, *24* (3), 461–483.

— and Oliver D. Hart, "Disclosure Laws and Takeover Bids," *Journal of Finance*, 1980, *35* (2), 323–334.

Grubb, Michael D., "Developing a Reputation for Reticence," *Journal of Economics Management Strategy*, 2011, *20* (1), 225–268.

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

Hennessy, Christopher A and Toni M Whited, "Debt dynamics," *The Journal of Finance*, 2005, *60* (3), 1129–1165.

— and — , "How costly is external financing? Evidence from a structural estimation," *The Journal of Finance*, 2007, *62* (4), 1705–1745.

Jovanovic, Boyan, "Truthful Disclosure of Information," *Bell Journal of Economics*, 1982, *13* (1), 36–44.

Judd, Kenneth L, *Numerical Methods in Economics*, MIT press, 1998.

Jung, Woon-Oh and Young K. Kwon, "Disclosure When the Market Is Unsure of Information Endowment of Managers," *Journal of Accounting Research*, 1988, *26* (1), 146–153.

Kadan, Ohad and Asaf Manela, "Estimating the value of information," *The Review of Financial Studies*, 2018, *32* (3), 951–991.

Kothari, Sabino P, Susan Shu, and Peter D Wysocki, "Do managers withhold bad news?," *Journal of Accounting Research*, 2009, *47* (1), 241–276.

Manela, Asaf, "The value of diffusing information," *Journal of Financial Economics*, 2014, *111* (1), 181–199.

Marinovic, Iván, "Dynamics of Financial Reporting," 2012. Rand Journal of Economics, forth.

Marinovic, Iván and Felipe Varas, "CEO Horizon, Optimal Pay Duration and the Escalation of Short-termism," Technical Report, Working Paper 2017.

_ , Andrzej Skrzypacz, and Felipe Varas, "Dynamic Certification and Reputation for Quality," 2015. Working paper.

Milgrom, Paul R., "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 1981, *12* (2), 380–391.

Shavell, Steven, "Acquisition and Disclosure of Information Prior to Sale," *Rand Journal of Economics*, 1994, *25* (1), 20–36.

Shin, Hyun Song, "News management and the value of firms," *Rand Journal of Economics*, 1994, pp. 58–71.

Tauchen, George, "Finite state markov-chain approximations to univariate and vector autoregressions," *Economics letters*, 1986, *20* (2), 177–181.

Taylor, Lucian A, "CEO wage dynamics: Evidence from a learning model," *Journal of Financial Economics, forthcoming. University of Pennsylvania working paper*, 2010.

Terry, Stephen J, "The Macro Impact of Short-Termism," 2016. Working paper.

Terry, Stephen, Toni M Whited, and Anastasia A Zakolyukina, "Information versus Investment," *Available at SSRN 3073956*, 2018.

Verrecchia, Robert E., "Discretionary Disclosure," *Journal of Accounting and Economics*, 1983, *5*, 179–194.

Zakolyukina, Anastasia A, "Measuring Intentional GAAP Violations: A Structural Approach," *Chicago Booth Research Paper*, 2014, (13-45).

Zhou, Frank, "Disclosure dynamics and investor learning," *Available at SSRN 2916276*, 2017.

A Appendix

A.1 Figures and Tables

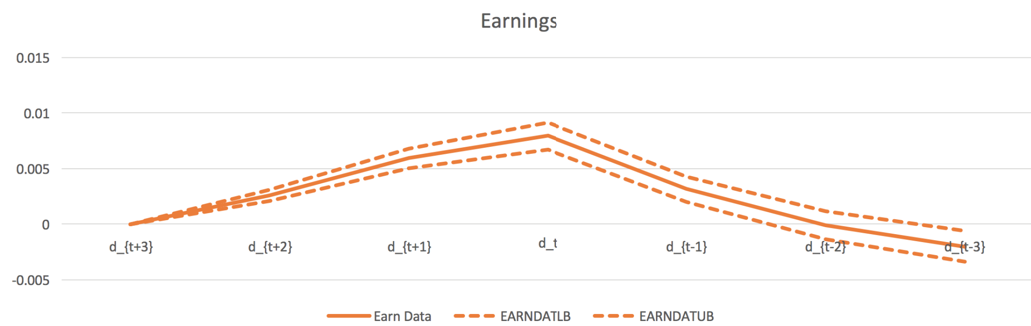

Figure 1. Earnings and Disclosure Histories

Note: This plot reports average realized earnings calculated among firms that forecasted the same earnings period d_t and average earnings k years before (after) the forecasted earnings period d_{t+k} (d_{t-k}), relative to earnings measured at $t-3$. A 95% confidence interval was obtained by bootstrapping 100 repetitions of each conditional mean.

Figure 2. Timeline

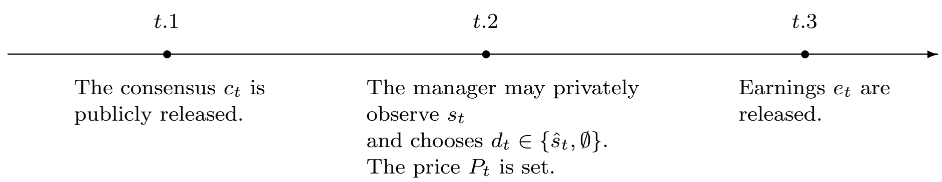

Table 1. Variable Definitions

Variables	Definitions
Descriptives	
Forecast frequency	Disclosure frequency by unique firm.
Total assets A_t	Total assets at period year end.
Management forecast	Raw one-year ahead management forecast scaled by total assets.
Realized IBES earnings	Raw realized one-year ahead earnings scaled by total asset.
IBES consensus	Consensus supplied by IBES prior to forecast or, if missing, median of 5 latest forecast by unique analysts or, if missing, lagged earnings.
Forecast surprise	Difference between management forecast and IBES consensus.
Forecast error	Difference between management forecast and realized IBES earnings.
Market capitalization	Price per share multiplied by fully-diluted number of shares.
Book to market	Ratio of year-end equity to market capitalization.
Model Variables	
e_t	Residuals of a regression of realized IBES earnings on firm and year fixed effects.
c_t	Residuals of a regression of realized IBES consensus on firm and year
s_t	Private signal of the manager.
θ_t	Indicator equal to one when manager is informed.
$break$ (br.)	Indicator equal to one when the firm changes disclosure at least once.
$dstart_t$	Indicator equal to one when a firm switches to non-disclosure.
$dend_t$	Indicator equal to one when a firm switches to disclosure.
Parameters	
λ_0 (resp., λ_1)	Prob. that the manager becomes uninformed (resp., informed) when when informed (resp. uninformed) in prior period.
σ_w	Standard deviation of manager signal.
ξ	Probability that the firm never discloses.
ζ	Probability that the firm is non-strategic and always discloses when informed.
σ_u	Standard deviation of earnings.
σ_v	Standard deviation of IBES consensus.
ξ	Probability that the firm never discloses.
ζ	Probability that the firm is non-strategic and always discloses when informed.
ρ	Earnings persistence.
β	Discount factor of the manager.
r	Discount rate on the firm's cash flow.

Table 2. Sample Selection

	Nb. of EA	Unique firms	Nb. of MF
I/B/E/S EA sample 2004-2016	83,935	10,542	
Matched to Compustat and CRSP,	67,145	9,665	
I/B/E/S CIG sample 2004-2016			
Matched to I/B/E/S EA	67,145	9,665	66,602
After prior EA date but prior period end	67,145	9,665	59,285
Minimum 6 month before period end	67,145	9,665	29,671
Retain only earliest MF	67,145	9,665	13,513
Must have market expectation	53,167	8,109	13,309
Full sample	**53,167**	**8,109**	**13,309**

Note: This table summarizes the sample selection criteria. Annual earnings announcements (EA) and management forecasts (MF) are obtained from I/B/E/S. Firms in our sample must be present in the CRSP and Compustat database using the I/B/E/S ticker, gvkey and permno matching tables from Wharton Research Data Services. Market expectation is calculated from the I/B/E/S analyst forecast file based on the median of the last five unique analyst forecasts.

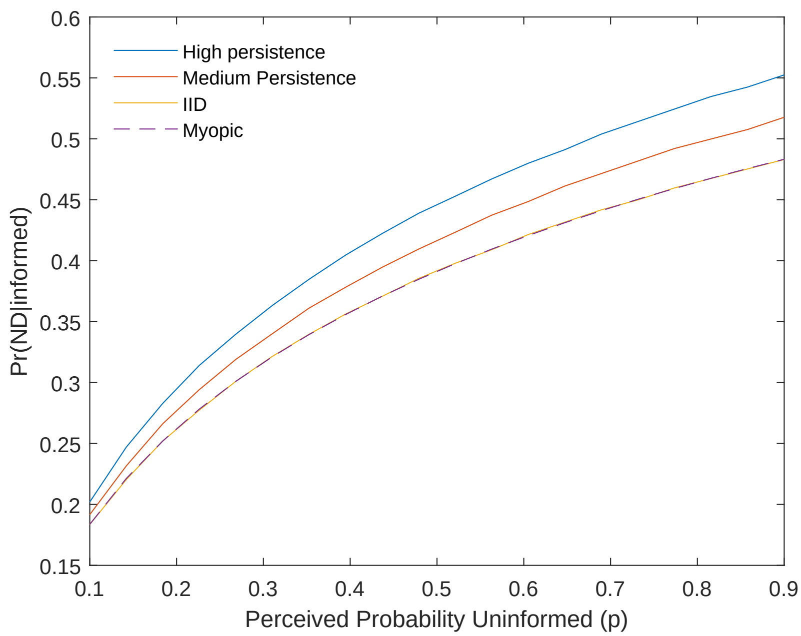

Figure 3. Effect of information persistence ($\overline{m}$) on disclosure choices

Note: The figure reports the likelihood of strategic withholding, i.e. the probability of non-disclosure $d = 0$ given an informed manager with $\theta = 1$, conditional upon a pre-existing level of market belief p that the manager is uninformed. The results were numerically computed from the stationary distribution of the model.

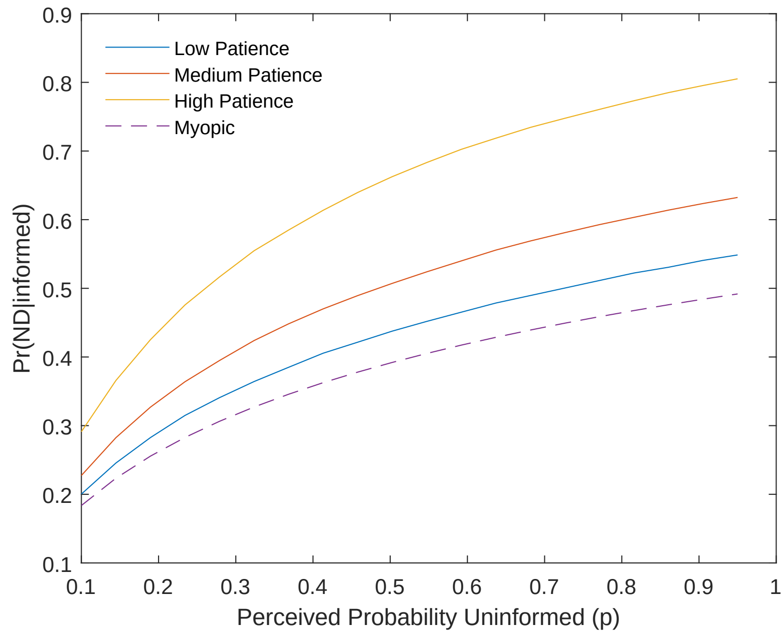

Figure 4. Effect of patience (β) on disclosure choices

Note: The figure plots the likelihood of strategic withholding, i.e. the probability of non-disclosure $d = 0$ given an informed manager with $\theta = 1$, conditional upon a pre-existing level of market belief p that the manager is uninformed. The results in solid lines were numerically computed from the stationary distribution of the model with various levels of manager patience β. The myopic model (dotted line) features disclosure policies equivalent to a model with an iid information endowment and independent of the level of manager patience β.

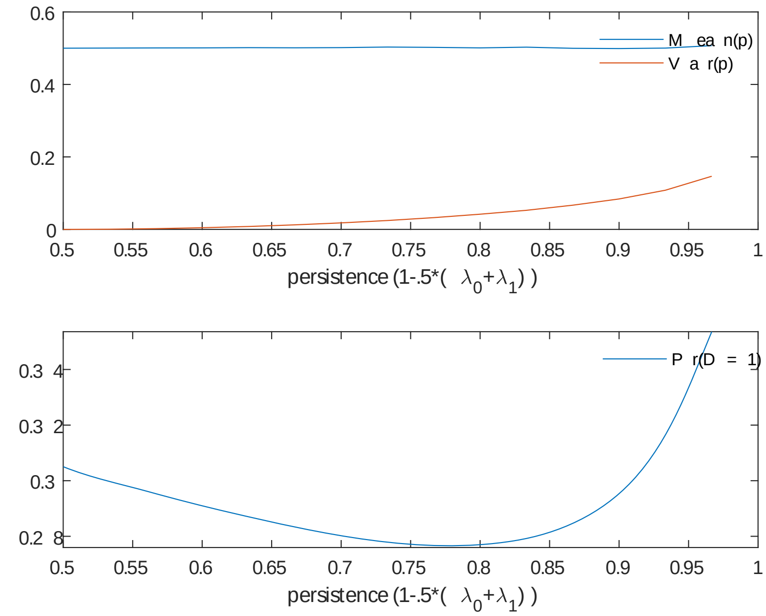

Figure 5. Effect of information persistence (r) on beliefs (p) and disclosure (d).
Note: The top panel of the figure plots the mean and variance of market beliefs p that the manager is uninformed in the steady state distribution, and the bottom panel computes the average likelihood of manager disclosure in the model. Both panels vary the level of persistence of the manager's information endowment $\overline{m} = 1 - \frac{\lambda_0+\lambda_1}{2}$ along the horizontal axis, holding $\overline{p}$ fixed.

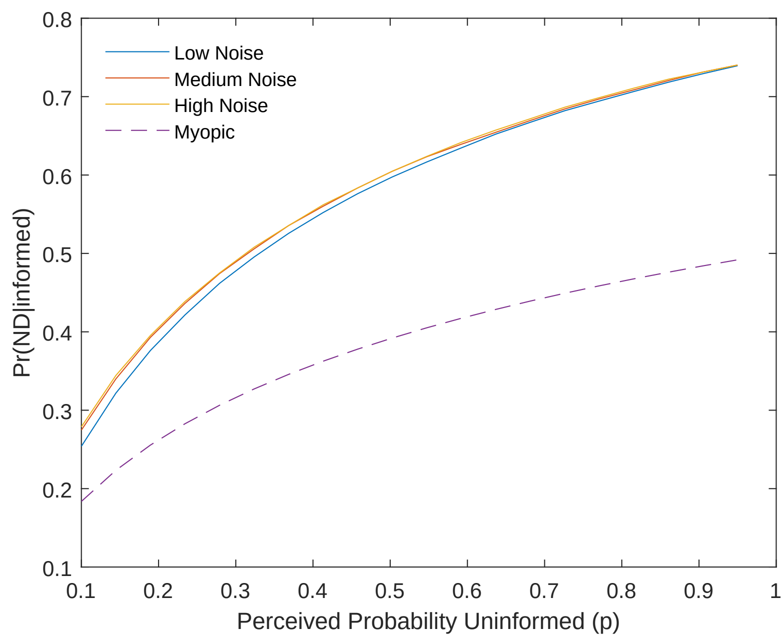

Figure 6. Effect of fundamental volatility (σ_u) on disclosure choices
Note: The figure plots the likelihood of strategic withholding, i.e. the probability of non-disclosure $d = 0$ given an informed manager with $\theta = 1$, conditional upon a pre-existing level of market belief p that the manager is uninformed. The results in solid lines were numerically computed from the stationary distribution of the model with various levels of fundamental volatility σ_u in earnings. The myopic model (dotted line) features disclosure policies equivalent to a model with an iid information endowment.

Table 3. Descriptive Statistics

	Mean	Std. Dev.	Min.	5%	25%	50%	75%	95%	Max.
Forecast characteristics									
Forecast frequency	23.49%	36.02%	0%	0%	0%	0%	39.23%	92.31%	100%
Management forecast	6.87%	4.39%	-10.68%	1.10%	3.80%	6.21%	9.23%	15.10%	23.81%
Forecast surprise	.000	.015	(.158)	(.014)	(.003)	(.001)	.001	.008	.523
Forecast error	.003	.031	(.256)	(.023)	(.006)	.001	.005	.038	.690
Realized I/B/E/S earnings	3.12%	9.92%	(71.6%)	(12.07%)	.91%	3.90%	7.67%	14.83%	25%
Firm characteristics									
Number of years	9.30	3.30	1	5	6	10	13	13	13
Total assets (bil.)	19.98	130.34	.003	.064	.374	1.381	5.222	50.096	3,065.553
Market capitalization (bil.)	7.782	26.038	.001	.051	.285	1.036	3.957	34,813	638,976
Book to market	.630	2.333	-337	.098	.300	.510	.809	1.626	31.73

Note: This table reports descriptive statistics. Forecast frequency is computed as the average of the frequency of management forecasts by firm. Forecast frequency (breaks) reports forecast frequency for firms with at least one change in disclosure choice. Management forecast (MF) is the management forecast scaled by assets where, in the case of a range forecast, we select the mid-point between each bound of the range. Realized I/B/E/S earnings are the pro-forma earnings reported by I/B/E/S. Forecast surprise is the difference between the MF and the market expectation (from I/B/E/S or estimated as lagged EPS when a consensus is unavailable). Forecast error is the difference between MF and the realized I/B/E/S earnings. Firm characteristics are obtained from Compustat. Market capitalization is obtained from CRSP and measured as the closing price multiplied by the fully dilluted number of shares. Book-to-market is the ratio of the firm's equity to its market capitalization.

Table 4. Outside Calibration of Earnings and Forecast Parameters

Parameters	Value	SE	Moment	Value	SE
$\mathbb{E}(e_t)$ (earnings)	.0120	.1431	$\mathbb{E}(c_t)$ (analyst)	.0277	.1265
ρ	.3055	.0289	$\text{Corr}(e_t, e_{t-1})$	.0653	.0192
σ_u (earnings)	.0749	.0018	$\text{St Dev}(e_t)$	.0682	.0014
σ_v (analyst)	.1383	.0098	$\text{St Dev}(c_t - e_t)$	.0658	.0010

Note: The first two columns report the name, role, and value of each of the three parameters calibrated externally through a moment-matching exercise. The final three columns report the data and model values of the three targeted moments. "Corr" refers to correlation and "St Dev" refers to standard deviation. In the model, moments are computed from a simulation of earnings e_t and consensus c_t processes for $10,000$ years, discarding the first 500 years of data. The minimization of the distance between model and data moments was performed numerically via particle swarm optimization. The number of firms is $5,056$ and the total number of observations is $31,883$.

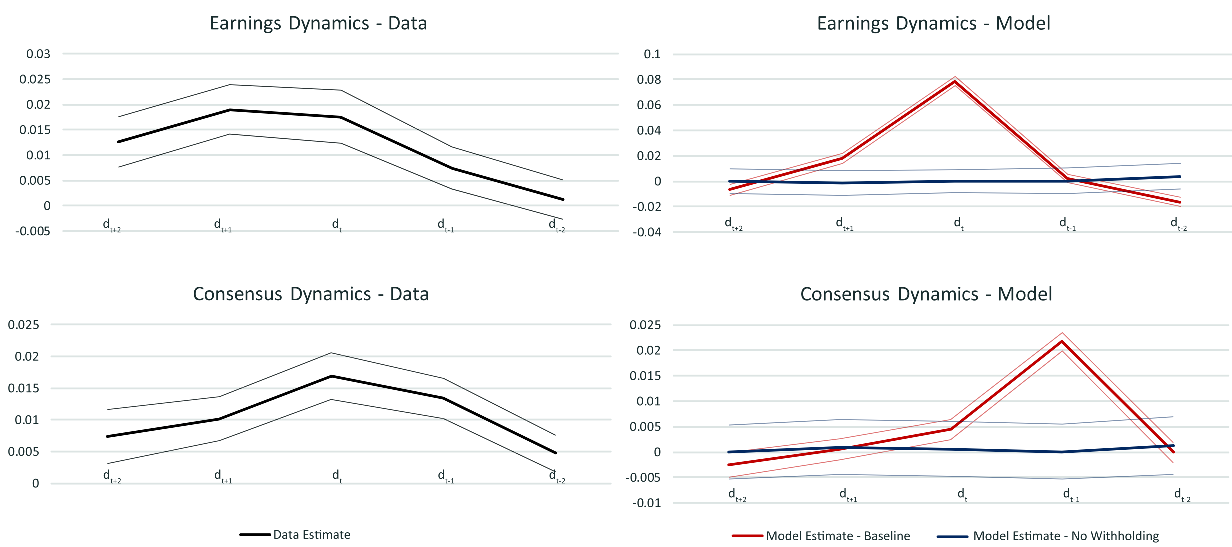

Figure 7. Earnings dynamics

Note: The top panel reports average earnings versus disclosure in the sample and model (with two standard errors as a dotted line).

Table 5. Targeted Moments for the SMM Estimation

Moment	Data Value	Baseline Model	No Withholding
P($d_t = 1$| break)	.4278 (.0102)	.5403 (.0334)	.4645 (.0135)
P($d_t = 1 \mid d_{t-1} = 1$, break)	.7820 (.0069)	.7622 (.0209)	.8102 (.0061)
P($d_t = 1 \mid d_{t-2} = 1$, break)	.6836 (.0092)	.6473 (.0270)	.6682 (.0098)
P($d_t = 1 \mid d_{t-3} = 1$, break)	.6152 (.0106)	.5885 (.0286)	.5630 (.0124)
St Dev($s_t - e_t \mid d_t = 1$, break)	.0215 (.0006)	.0196 (.0009)	.0222 (.0007)
P($d_t = 1$ Never)	.6169 (.0071)	.6424 (.0191)	.6286 (.0195)
P($d_t = 1$ Always)	.0789 (.0039)	.0831 (.0171)	.1269 (.0140)
Number firms	5,056		
Number total	31,883		

Note: The table reports the data values of moments computed from I/B/E/S data on manager forecasts s_t conditional upon disclosure d_t, and the values of the targeted moments in the estimated model. No (strategic) withholding restricts the model so that firms set $d_t = \theta_t$ with prob. $1 - \xi$. The first five moments are calculated conditional on firms with variation in disclosure. The first Standard errors are obtained by block bootstrap, resampling at the firm level.

Table 6. SMM Parameter Estimates

Parameter, Role	Estimate	No withholding
λ_0, $\mathbb{P}(\theta_t = 0 \mid \theta_{t-1} = 1)$	.0367 (.0068)	.1004 (.0099)
λ_1, $\mathbb{P}(\theta_t = 1 \mid \theta_{t-1} = 0)$	.0561 (.0108)	.0395 (.0055)
σ_w, St Dev Manager Signal	.0154 (.0008)	.0178 (.0008)
$\xi, \mathbb{P}$ (Never d_t)	.4339 (.0637)	.1454 (.0881)
$\zeta, \mathbb{P}$ (Always $d_t = \theta_t$)	.0385 (.0062)	

Note: Baseline parameters estimated via SMM, targeting the five moments described in the text and implemented with a diagonal weighting matrix and numerical minimization. No (strategic) withholding restricts the model so that firms set $d_t = \theta_t$ with prob. $1 - \xi$. Standard errors are in parenthesis. The model-implied moments come from a simulation of $20,000$ firms for 9 years each, a symmetric structure to the data. The standard errors are based on an empirical moment covariance matrix computed from a block bootstrap procedure, resampling at the firm level, together with a moment Jacobian computed from numerical forward differentiation averaging over step sizes of .25%, .5%, and 1%.

Table 7. Disclosure and selection

Panel A: Disclosure and Selection			
	Data	Model	Model
$\frac{e_t}{A_t}$		Baseline	No withholding
d_t	.018***	.056***	.000***
	(.003)	(.002)	(.002)
Fixed Effects	Year, Firm	Year, Firm	Year, Firm
Years	2004-2016	2004-2016	2004-2016
Firms	5,023	5,000	5,000
Firm-Fiscal Yr. Obs.	31,246	34,781	34,754
Panel B: Disclosure breaks			
Dependent var.	Data	Model	Model
$\frac{e_t}{A_t} - \frac{e_{t-1}}{A_{t-1}}$		Baseline	No withholding
$dstart_t$	.004*	.056***	.000
	(.002)	(.002)	(.000)
$dend_t$	-0.006**	-0.065***	.000
	(.003)	(.002)	(.000)
d_t	.002***	.000	.000
	(.001)	(.001)	(.000)
Fixed Effects	Year, Ind.	Year, Firm	Year, Firm
Years	2004-2016	2004-2016	2004-2016
Firms	4,990	5,000	5,000
Firm-Fiscal Yr. Obs.	30,906	34,781	34,754
Panel c: Disclosure spells			
Dependent var.	Data	Model	Model
d_t		Baseline	No withholding
e_t/A_t	.082***	.581***	-0.007
	(.017)	(.029)	(.016)
d_{t-1}	.267***	.266***	.537***
	(.014)	(.011)	(.008)
$e_t/A_t \times d_{t-1}$	.639***	1.951***	-0.037
	(.110)	(.085)	(.052)
Fixed Effects	Year, Ind.	Year, Firm	Year, Firm
Years	2004-2016	2004-2016	2004-2016
Firms	5,023	5,000	5,000
Firm-Fiscal Yr. Obs.	31,246	34,781	34,754

Note: Panel A reports the results of an ordinary least-square regression of earnings on disclosure. Panel B reports the results of an ordinary least-square regression of change in earnings, on disclosure breaks (break) defined as forecasting changes between consecutive periods, and an interaction. Standard errors are in parenthesis. Panel C reports the results of an ordinary least-square regression of forecasts on earnings , lagged disclosure defined as an indicator variable equal to one when a disclosure was made in the prior period, and an interaction between these two variables. Firm and year fixed effects are included and earnings are scaled by current assets.

Table 8. Counterfactual Forecast Errors

	RMSE	relative RMSE	$\mathbb{P}(ND\|Informed)$
No information	7.5%	100%	
Benchmark (a)	6.3%	84%	.26
No withholding (b)	6.0%	80%	.00
Noisy forecasts	6.5%	87%	.48
No consensus	6.6%	88%	.21
ΔRMSE (a)-(b)	0.4%		
Std. error	(0.01%)		

Note: The Baseline Model results are obtained by simulating a panel of 1,000 firms for 5,000 periods on the basis of the estimated annual parameters. No Information model runs the simulation with very high noise on consensus and management signal. No Withholding runs the simulation by requiring the manager to always disclose when informed. Noisy forecasts (No consensus) runs the simulation with very high noise on the manager signal (on the consensus). Std. error is calculated by bootstrapping Benchmark and No Withholding with model parameters sampled from the estimated parameter covariance matrix.

Table 9. Robustness Analyses

		relative RMSE	relative RMSE (no withholding)	$\mathbb{P}(ND\|Informed)$
Baseline		84%	54%	.26
High	$\rho = .367$	84%	81%	.27
Low	$\rho = .291$	84%	81%	.26
High	$\sigma_u = .079$	87%	84%	.27
Low	$\sigma_u = .068$	77%	74%	.26
High	$\sigma_w = .166$	84%	81%	.24
Low	$\sigma_w = .115$	82%	80%	.29
High	$\beta = .833$	84%	80%	.28
Low	$\beta = .663$	83%	80%	.27

Note: The Baseline Model results are obtained by simulating a panel of 1,000 firms for 5,000 periods on the basis of the estimated annual parameters, and the parameter in column 1. Rows 2-3 report in earnings correlation, rows 4-5 report changes in fundamental earnings noise, rows 6-7 report changes in manager signal noise and rows 8-9 report changes in manager impatience.

Table 10. Subsample Analysis

Subsample Estimates	High nb. of analysts	Low nb. of analysts	High size	Low size	High cons. errors	Low cons. errors
λ_0, $\mathbb{P}(\theta_t = 0\|\theta_{t-1} = 1)$	.05 (.01)	.10 (.01)	.06 (.01)	.06 (.02)	.05 (.00)	.06 (.01)
λ_1, $\mathbb{P}(\theta_t = 1\|\theta_{t-1} = 0)$	.17 (.02)	.03 (.01)	.20 (.03)	.09 (.04)	.12 (.01)	.19 (.04)
σ_w, std manager signal	.25 (.05)	.02 (.00)	.11 (.01)	.04 (.00)	.04 (.01)	.22 (.04)
$\xi, \mathbb{P}$ (Never d_t)	.50 (.02)	.44 (.09)	.36 (.03)	.52 (.09)	.59 (.02)	.61 (.02)
$\zeta, \mathbb{P}$ (Always d_t)	.13 (.02)	.01 (.05)	.12 (.05)	.04 (.03)	.00 (.00)	.07 (.06)
RMSE	58%	85%	70%	68%	77%	81%
RMSE, no concealment	58%	81%	69%	61%	68%	81%
$\mathbb{P}$(ND|Informed)	.24	.38	.27	.30	.26	.25
Nb of firms	2,507	2,549	2,230	2,826	2,626	2,432
Nb of obs.	17,396	14,487	17,178	14,705	15,574	16,309
Subsample Moments: data (se) model						
E($d_t = 1$ | br.)	.47 (.01) .59	.34 (.01) .40	.46 (.01) .57	.39 (.01) .50	.42 (.01) .53	.43 (.01) .58
P($d_t = 1$ | $d_{t-1} = 1$, br.)	.81 (.01) .75	.72 (.02) .64	.81 (.01) .73	.74 (.01) .71	.77 (.01) .72	.79 (.01) .75
P($d_t = 1$ | $d_{t-2} = 1$, br.)	.72 (.01) .67	.59 (.02) .51	.72 (.01) .64	.63 (.01) .59	.67 (.01) .62	.70 (.01) .66
P($d_t = 1$ | $d_{t-3} = 1$, br.)	.65 (.01) .62	.51 (.02) .44	.65 (.01) .60	.56 (.01) .53	.59 (.01) .56	.63 (.01) .61
Std ($s_t - e_t$ | $d_t = 1$, br.)	.02 (.00) .04	.03 (.00) .02	.01 (.00) .02	.03 (.00) .04	.02 (.00) .04	.02 (.00) .03
P($d_t = 1$ Never)	.50 (.01) .56	.73 (.01) .84	.54 (.01) .43	.67 (.01) .67	.62 (.01) .67	.61 (.01) .65
P($d_t = 1$ Always)	.13 (.01) .13	.03 (.00) .01	.13 (.01) .12	.04 (.00) .04	.07 (.01) .04	.09 (.01) .09

Parameters estimated via SMM as in Table 10 described in text. Breaks (br.) indicate a firm changing disclosure behavior at least once. Number of analysts is the number of unique analysts having made a forecast from the prior earnings announcement to the current earnings announcement; size is the market capitalization and consensus errors is the standard-error between the analyst consensus and the realized earnings. We form each group by averaging by firm and selecting the subsample depending on whether the firm is above (high) or below (low) the full sample average.

A.2 Proofs

Proof. [Proof of Proposition 2] This result comes from noting that when $\lambda_0 = 1 - \lambda_1$, we get

$$V^D(p, z, s) - V^{ND}(z, s) = P^D(p, z, s) - P^{ND}(p, z). \tag{25}$$

which means that the manager maximizes his myopic price gain when choosing whether or not to disclose. □

Proof. [Proof of Proposition 3] By assumption, the value function can be expressed as ,where the last equality follows the law of iterated expectations and AR(1) property of earnings. □

Proof. [Proof of Proposition 4] For expositional purposes, we normalize the unconditional mean of e to zero because it plays no role in the proof, and we drop the dependence on z.

(1) Suppose that there is an equilibrium, such that information is withheld (case a) or

disclosed (case b) for all s conditional on some belief $p \in (0,1)$,

$$\begin{aligned} V_1^{ND}(p,s) &= P^{ND}(p) + \beta E[\lambda_0 V_0(p') + (1-\lambda_0)V_1(p',s')|s] \\ &= \beta E[\lambda_0 V_0(p') + (1-\lambda_0)V_1(p',s')], \end{aligned}$$

where $p' = p(1-\lambda_1) + (1-p)\lambda_0$ (case a) or $p' = 1-\lambda_1$ (case b) so that $V_1^{ND}(p,s)$ does not depend on s. On the other hand, from proposition 3, $V_1^D(p,s)$ is unbounded in s, implying that the manager would prefer to disclose for s large enough (case a) or withhold for s small enough (case b), which is a contradiction.

(2) Suppose that there is an equilibrium, such that the manager adopts a threshold strategy k_p for any p, i.e., discloses if $s > k_p$ and withholds if $s < k_p$. We know from (1) that k_p is finite. Then,

$$V_1^{ND}(p,s) = P^{ND}(p) + \beta E[\lambda_0 V_0(p') + (1-\lambda_0)V_1(p',s')|s], \tag{26}$$

where $p' = p(1-\lambda_1) + (1-p)\lambda_0$ if $s < k_p$ and $p' = 1-\lambda_1$ if $s > k_p$.

(2.a) Suppose that

$$E(\lambda_0 V_0(p(1-\lambda_1) + (1-p)\lambda_0) + (1-\lambda_0)V_1(p(1-\lambda_1) + (1-p)\lambda_0, s')) \tag{27}$$

$$= E(\lambda_0 V_0(1-\lambda_1) + (1-\lambda_0)V_1(1-\lambda_1, s')) = 0, \tag{28}$$

where the last equality follows from proposition 2. It follows that $E(\lambda_0 V_0(p') + (1-\lambda_0)V_1(p') = 0$ for any belief p' that may occur on the equilibrium path, and k_p must be the static threshold in proposition 2. Then, $V_0(p') = P^{ND}(p')$, which is increasing in p', and $V_1(p',s)$ is not a function of p' from proposition 3, which is a contradiction to (28).

(2.b) Suppose that

$$E(\lambda_0 V_0(p(1-\lambda_1) + (1-p)\lambda_0) + (1-\lambda_0)V_1(p(1-\lambda_1) + (1-p)\lambda_0, s')) \tag{29}$$

$$< E(\lambda_0 V_0(1-\lambda_1) + (1-\lambda_0)V_1(1-\lambda_1, s')) = 0, \tag{30}$$

so that $E(\lambda_0 V_0(p') + (1-\lambda_0)V_1(p',s')) < 0$ for any p' on the equilibrium path. Because $\mathbb{E}(V_1(p',s')) \geq V_0(p')$ for any p', it must hold that

$$E(p' V_0(p') + (1-p')V_1(p',s')) \leq E(\lambda_0 V_0(p') + (1-\lambda_0)V_1(p',s')) < 0, \tag{31}$$

for any p' on the equilibrium path. The left-hand side of this inequality is equal to the unconditional mean of e from the law of iterated expectations, which is a contradiction.

(2.c) Suppose that

$$E(\lambda_0 V_0(p(1-\lambda_1) + (1-p)\lambda_0) + (1-\lambda_0)V_1(p(1-\lambda_1) + (1-p)\lambda_0, s')) \tag{32}$$

$$> E(\lambda_0 V_0(1-\lambda_1) + (1-\lambda_0)V_1(1-\lambda_1, s')) = 0. \tag{33}$$

Note that $V_1^D(s)$ is increasing and continuous in s from proposition 3 so that we must have

$V_1^{ND}(p, k_p - a) \leq V_1^D(k_p) \leq V_1^{ND}(p, k_p + a)$ for any $a > 0$, where equation (33) implies that at least one inequality is strict. Assume that $V_1^{ND}(p, k_p - a) < V_1^D(k_p)$. Then, there exists $\epsilon > 0$ sufficiently small, such that $V_1^{ND}(p, k_p - a) < V_1^D(k_p - \epsilon)$, which contradicts that k_p is the disclosure threshold. The case of $V_1^{ND}(p, k_p - a) > V_1^D(k_p)$ follows from a symmetric argument. □

A.3 Solving the Model

The computational strategy we use is policy iteration. The steps are as follows:

1. Discretize the state space.

2. On the s-th iteration of the solution algorithm, guess a disclosure policy $d^{(s)}(p, e_{-1}, c, s)$.

    (a) Assume that market beliefs and manager actions are governed by $d^{(s)}$, and iterate forward on the system of Bellman equations above until the implied $V_1^{(s)}$, $V_0^{(s)}$ converge to some tolerance.

    (b) Compute the stationary distribution $\mu_1^{(s)}(p, e_{-1}, c, s)$ and $\mu_0^{(s)}(p, e_{-1}, c)$ of the model given $d^{(s)}$, as well as the exogenous distributions in the model. This involves repeatedly pushing forward weight on a histogram given the policies and exogenous transitions until the distributions stabilize to within some tolerance.

    (c) Compute a *new* policy $d^{(s+1)}(p, e_{-1}, c, s)$, simply given by

    $$\arg\max_d \left(dV_1^{D(s)} + (1 - d)V_1^{ND(s)} \right). \tag{34}$$

    (d) Then, compute an error measure given by the mean absolute difference between $d^{(s+1)}$ and $d^{(s)}$, weighted by the ergodic distributions $\mu_1^{(s)}$ and $\mu_0^{(s)}$. This error is exactly equal to the probability of disclosure policy deviation given assumed market beliefs. When this error is sufficiently small, you have computed an equilibrium.

3. Once we have solved the model, we can compute moments as desired for input into the structural estimation routine.

We implement our solution algorithm in Fortran, discretizing driving exogenous processes for earnings, consensus forecasts, and manager signals using the method of Tauchen (1986). Broadly, our numerical approach to the resulting discrete-state dynamic programming problem follows the methods outlined in Judd (1998).