Disclosing a Random Walk*

Ilan Kremer[†] Amnon Schreiber [‡] Andrzej Skrzypacz[§]

September 23, 2021

Abstract

We examine a dynamic disclosure model in which the value of an asset follows a random walk. Every period, with some probability an agent learns the value of the asset and decides whether to disclose it. The agent maximizes the market perception of the asset's value, which is based on disclosed information. We show that in equilibrium the agent follows a threshold strategy but also reveals pessimistic information that reduces market perception. We examine different variants and show, for example, that when he can disclose stale information he is less likely to disclose current values, but over time discloses more.

JEL Codes: D21, D82, D83, G12, L11.

1 Introduction

Analyzing voluntary disclosure and its consequences for the functioning of markets intersects accounting, economics, and finance. The existing liter-

*We thank seminar participants, our colleagues, and Daniel Chen and Suraj Malladi for comments and suggestions.

[†]The Hebrew University and University of Warwick. Email: ikremer@huji.ac.il.

[‡]Bar-Ilan University. Email: amnonschr@gmail.com

[§]Stanford Graduate School of Business, Stanford University. Email: skrz@stanford.edu

ature on strategic disclosure focuses on static or dynamic models with the asset having a fixed value. This is at odds with a common feature of financial and other markets in which, as information arrives, the (expected) value of an asset follows a random walk. Such stochastic evolution introduces new strategic considerations. The agent may decide to hide information, hoping to disclose the value in the future when it is higher. However, when the agent continues to hide information, the market becomes more pessimistic about the value. Therefore, in deciding whether to disclose information, the agent accounts for the effect of current and future market perceptions. Optimal disclosure decisions depend on his expectations about the evolution of the future value of the firm and the evolution of the market's beliefs. Our goal in this paper is to provide a first step in bridging the gap between static and dynamic voluntary disclosure models.

We study a model of dynamic voluntary information disclosure by a manager of a public firm. The value of the firm follows a random walk. In every period, with some probability, the manager holds material information and chooses whether to disclose it. A key feature of strategic disclosure models captures the fact that it is often difficult to blame an agent for not disclosing information, as it is hard to prove whether he had it. However, he can be punished if he discloses false information. We follow this logic and assume that the agent may disclose information only when it is timely. A regulator may not know with certainty that an agent has concealed information, so it cannot punish him for this. However, when the agent has delayed revealing information, it is clear the agent has concealed information so he can be punished. [1] We assume that the market sets the current price as the expected value of the firm conditional on the public history of disclosures. The equilibrium is based on the manager maximizing a weighted average of market prices, and market prices being consistent with the manager's strategy. In

[1] To understand how this restriction affects the equilibrium outcome, in Section 5.1 we allow the agent to disclose stale information.

particular, the market no-disclosure price is based on the set of values the manager chooses not to disclose.

We first show that similarly to static models, the equilibrium is based on threshold strategies. At any given time, and given some history, the manager who has information reveals it if and only if it exceeds a certain history-dependent threshold. In a dynamic environment, this feature has to some interesting implications. The value process is assumed to be a martingale. Prices also follow a martingale as they reflect expected values conditional on public information. However, since the agent follows a threshold strategy, prices are more positively skewed, compared to the value process. In equilibrium, "no news is bad news": in the intervals between disclosures, prices drift down and there is usually (but not always; see our second result) a positive jump upon disclosure.

We then characterize the equilibrium thresholds for disclosure. We show that disclosure thresholds are always lower than the no-disclosure price (apart from the last period). This implies that with positive probability, the agent discloses information that leads to a lower price than the price that would have prevailed had he decided not to disclose. This result stands in contrast to a one-period model or a myopic behavior. The intuition for why the manager chooses to disclose some values even though by doing so he reduces prices today is that by disclosing today, he reduces future uncertainty. It is important to note that the manager is risk-neutral, and prices are equal to the expected value of the firm regardless of what uncertainty the market faces. So, it is perhaps surprising that despite being risk-neutral, the manager discloses information that reduces uncertainty about the value of the firm at the cost of reducing the current price. The key reason for this behavior is the difference in beliefs about future values between the manager and the market: the market forms beliefs based on the public disclosure history. The manager additionally knows the undisclosed information. These differences are generated over time and are affected by the disclosure decisions of the manager. A decision to withhold information leads to higher uncertainty (i.e., the market's beliefs being more dispersed than the agent's). This higher uncertainty implies lower no-disclosure prices in the future, as the market accounts for a fatter left tail. The market is skeptical about the value since there is a chance the agent is hiding information, and the skepticism is higher the more uncertain the value is from the market's point of view.

We then establish a generalization of the "minimum principle" that was introduced in Acharya et al. (2011), by defining the "suspicious belief principle". If the market believes that the agent follows a certain disclosure strategy, then it sets non-disclosure prices as the expected values conditional on no-disclosure. We show that the equilibrium disclosure strategy must satisfy a certain pessimistic beliefs property (for computation of expected no disclosure prices) for all possible disclosure strategies. This provides a necessary and sufficient condition for a strategy to be an equilibrium strategy. It takes a simple form if we assume that the agent cares only about the final price. We rely on this solution in Section 5.1 when we consider the disclosure of stale information.

To better understand our model, we study a few variants in Section 5. We first consider the restriction that the agent must disclose timely information. Relaxing this restriction implies that the agent can disclose at time t the value $v_{t'}$ that he had obtained at an earlier date $t' < t$. We analyze a two-period case. With the option to disclose stale information, off-equilibrium beliefs play a more significant role, and there exists an equilibrium in which the agent never discloses stale information. However, when we apply a refinement that is based on a trembling hand, there exists a unique equilibrium outcome in which the manager does disclose stale information.

We show that the manager follows a myopic strategy in the sense that he discloses a piece of evidence (either current or stale) if and only if it leads to a higher current price. While this is similar to the static case, the disclosure strategy is more involved. For example, at $t = 2$ the disclosure strategy is defined by two thresholds, as he may disclose two different pieces of information. We then ask how the ability to disclose stale information affects overall equilibrium disclosure. The first-period equilibrium threshold in the stale model is higher than in our main model. However, the equilibrium threshold for disclosing the first-period value decreases in period two, so that the agent discloses with delay some values that he conceals in the first period. We prove that the threshold in the second period is even lower than the first-period threshold when only timely disclosure is allowed. The takeaway is that while our assumption that information must be disclosed in a timely manner may seem at first to increase disclosure, this outcome may reverse when we consider the aggregate/long-term effects.

We then examine a variant in which an information event occurs in each period with some probability. This information event leads to a change in the asset's value that can be disclosed. Since it is uncertain whether such an event has occurred, the manager has discretion whether to disclose it or not. We first note that if the disclosure is about the asset's value, the same results as in our main model continue to hold. However, when the information the manager can disclose is only about the changes in the asset's value, then the equilibrium is similar to a static model, and the disclosure strategy is myopic. This is because the changes in the asset's value are independent over time. This highlights that the serial auto-correlation in the asset's value leads to a non-myopic disclosure strategy.

1.1 Related Literature

The voluntary disclosure literature goes back to Grossman and Hart (1980), Grossman (1981), and Milgrom (1981), who showed that if it is commonly known that the agent is privately informed, then there is full disclosure. Our paper follows Dye (1985) and Jung and Kwon (1988), who showed in a one-period model that when investors are uncertain about the information endowment of the agent, then there exists a non-trivial equilibrium in which some types withhold information and some disclose it. As mentioned in the introduction, despite the vast literature on voluntary disclosure, only a few papers have examined multi-period settings in which the information changes over time. Shin (2003) and Shin (2006) study a setting in which a firm may learn a binary signal for each of its independent projects, and each project may fail or succeed. In this binary setting, Shin studies the "sanitization" strategy, under which the agent discloses only the good (success) news. The timing of disclosure does not play a role in such a setup. Pae (2005) considers a single-period setting in which the agent can learn up to two normally distributed signals.

Einhorn and Ziv (2008) study a setting in which in each period the manager may obtain a single signal about the period's cash flows. At the end of each period, the realized cash flows are publicly revealed, which eliminates the dynamic considerations that are at the heart of the present paper. Acharya et al. (2011) examine a dynamic model in which a manager learns one piece of information at some random time, and his decision to disclose it is affected by the release of some external news. Bertomeu et al. (2011) study a reputation model in which the manager may learn a single private signal in each of the two periods. The manager can be either "forthcoming" and disclose any information he learns, or he may be "strategic". At the end of each period, the firm's signal/cash flow for the period becomes public, and the market updates beliefs about the value of the firm and the type of the agent. Importantly, the option to "wait for a better signal" that is behind our main result is not present in any of these papers.

Guttman et al. (2014) examine a two-period model in which there are potentially two pieces of information. The main result is that in the second period, the market values the same signal more if it is disclosed in the second period rather than in the first period. This is different from the result we obtain for the case with stale information disclosure. The reason for this difference stems from the nature of the evidence we consider. In our model, the evidence is about the current value, and so it has a "time stamp". When it is disclosed late, the market knows that it was disclosed with delay. Moreover, the information sets in the different periods can be described as filtration. Information in the first period is not only less informative than information in the second period, but also a garbling of the value in the second period. Aghamolla and An (2019) study a two-period model in which the value is normally distributed. They obtain a result that is similar to our first result. The manager may disclose negative information in the first period. The setup and the driving force, however, are quite different. In their paper, it is commonly known that the manager knows the value of the firm. Disclosure, however, is costly. When deciding whether to disclose, the manager optimizes the benefits of disclosure against the cost of disclosure in each period.

2 The Model

We consider a model of dynamic strategic disclosure with a single agent who can be viewed as a manager of a public firm. Time is discrete, $t \in \{1, 2..., T\}$. The starting value of the firm, V_0, is known. The value evolves as a random walk

$$V_t = V_0 + \sum_{\tau=1}^{t} \Delta V_\tau,$$

with increments $\Delta V_\tau \equiv V_\tau - V_{\tau-1}$, which are zero-mean i.i.d. random variables, with cumulative distribution function F, strictly positive density function f, and finite variance.

In every period t, with probability $\pi \in (0, 1)$ the manager learns the current value of the firm, and can credibly/verifiably disclose it.[2]

[2]For simplicity, we assume that the event of being able to disclose is independent of the realized value. Our model and results can be generalized to allow the distribution of the increments to depend on whether the agent can disclose information, as long as we maintain that the current value and future opportunities to disclose are independent. See Section 5.2 for an example of such a model.

A strategy of the agent is a disclosure rule that specifies which values of V_t to disclose. Denote by $H_t = \{d_1...d_t\}$ the (public) history of disclosures at time t, where $d_\tau = V_\tau$ if the agent discloses the value of the firm, and $d_\tau = \varnothing$ if he does not. The manager can reveal the value of the firm only when he has verifiable information. The manager's time-t disclosure strategy is denoted by $\sigma_t(H_{t-1}, V_t) \in [0, 1]$, which is the probability of disclosure of the current value if he can disclose it.

The market sets prices to be equal to the expected value of V_t based on rational expectations and conditional on the history of the agent's disclosures. We denote the market prices by $P_t(H_t)$.

We show later that the equilibrium disclosure strategies are threshold rules. Thus, it is convenient to define the pricing function $\hat{P}_t(H_t, x_1, ..x_t)$ as the expectation of V_t conditional on the history and conditional on the manager using some arbitrary disclosure thresholds $x_1, .., x_t$ in periods when he did not disclose information.[3]

The manager maximizes a weighted sum of prices:

$$\sum_{t=1}^{T} w_t \cdot P_t(H_t),$$

for some known weights $w_t \geq 0$. This general specification allows us to capture a standard discounted utility model as well as the case in which managerial compensation is more sensitive to stock prices on specific dates.

Definition 1. *An equilibrium of this model is a disclosure strategy of the agent and market prices, $\{\sigma_t(H_{t-1}, V_t), P_t(H_t)\}$ such that:*

1. *(Sequential Rationality) After every history, the manager maximizes his expected utility given market prices.*

2. *(Sequential Consistency of Prices). The price at time t equals the expected value of V_t conditional on the public history and the manager's*

[3] For equilibrium thresholds it holds that $\hat{P}_t(H_t, x_1, ..x_t) = P_t(H_t)$.

disclosure strategy.

Since we presented the equilibrium using a verbal description rather than a mathematical expression, we should make a few comments. When the manager maximizes his utility given prices, he considers how disclosure affects his current price and also future prices. In computing prices, there are two cases to consider. In periods in which the agent discloses information, i.e., for any H_t such that $d_t = V_t$, we have $P_t(H_t) = V_t$ (since we assumed disclosure is credible). In periods in which the agent does not disclose, i.e., for any H_t such that $d_t = \varnothing$, let $\tau \in \{0, ...t-1\}$ be the last period in which $d_\tau \neq \varnothing$ (or 0 in case the agent had never disclosed). For those histories, $P_t(H_t) = E[V_t | V_\tau, \{\sigma_s(H_{s-1}, V_s), d_s = \varnothing\}_{s=\tau+1}^{s=t}]$. Note that the agent's strategy $\{\sigma_t(.,.)\}$ since the last disclosure is used to calculate the expected value conditional on no disclosure.

The equilibrium conditions apply to histories on and off the equilibrium path. In case the history is on the equilibrium path, prices follow Bayes' rule. The off-the-equilibrium path events are disclosures of values to which σ assigns zero probability. After those histories, when the agent discloses value V_τ, we require that the continuation strategy and the prices form an equilibrium of the model as if the starting value were $\widetilde{V}_0 = V_\tau$ and the model had horizon $\widetilde{T} = T - \tau$.[4]

Discussion of Assumptions We assumed that the increments ΔV_t have identical distributions and that the probability that the agent can disclose information, π, is constant. These assumptions are just for ease of exposition. Our qualitative results are robust to allowing F and π to vary over time. Our assumption that increments have zero mean and are independent implies that the value process follows a martingale. This is a common feature of

[4]It would also be possible to write the model as a game between the agent and a competitive set of investors. In that case, we could use as the equilibrium notion perfect Bayesian equilibrium with the requirement that beliefs satisfy proper subgame consistency (see Kreps 2020).

finance models that captures the fact that in efficient markets, prices are equal to expected value conditional on available information. Prices then follow a martingale, as the expected value conditional on an increasing set of information is a martingale process.

A more substantial assumption is that the agent may reveal only timely information. If he does not reveal V_t at t, he cannot reveal it later. One reason for this assumption is that the verifiable information may be short-lived. Another reason is that agents are commonly required to reveal all material information in a timely fashion. An agent can hide information by pretending an information event has not occurred. If he reveals the information with delay, it becomes clear he has violated that requirement. For example, an outside investor could approache the manager with a credible price offer (for the firm or a part of it). When the offer is made, the agent has to decide whether to treat it as material information. If he does, he is compelled to reveal it immediately to the shareholders. If he does not, the offer expires, and hence it may not be possible to reveal it credibly. Moreover, revealing it later could put the manager in trouble. We later examine the effect of this assumption on equilibrium disclosure by studying a case in which the manager can also disclose stale information (see Section 5).

We also assumed that the agent learns V_t only when he can disclose it (for example, news about V arrives via offers made by outside investors). Our analysis and results would remain unchanged if instead the agent observed V_t even in periods when he could not disclose it. This is because the agent can only take actions when he has verifiable information and once that information arrives, the current realization of V_t is a sufficient statistic for the firm's future value.

Finally, we assumed that when the agent obtains information, it fully reveals the value of the firm V_t. Instead, we could interpret V_t as the best (possibly noisy) estimate available to the agent at the time of disclosure. For example, if instead of observing at t the value V_t the agent observed the value of V_s for some $s < t$, then, as long as the time he observes the signal can be verified and the agent is constrained to reveal V_s only in the period he observes it, our model with timely disclosure would apply. In other words, what matters in the model for information to be "timely" is not whether the information is about current or past values of V_t, but rather when the agent received that information: if he can disclose the information only in the period in which he received it, our model still applies.

3 Preliminary Analysis

3.1 Disclosure in a One-Period Model

We begin the analysis by considering a static model (which is useful for understanding disclosure in period T). We review some results that are due to Dye (1985), Jung and Kwon (1988), and Acharya et al. (2011). These results form a benchmark against which we will compare disclosure policies in the dynamic model.

In the static model, the asset value is given by V_1 and is distributed according to some distribution F. With probability π, the manager learns the value and can decide whether or not to disclose it. In equilibrium, the manager follows a threshold strategy. The manager withholds the value if and only if it is below a certain threshold. The reason is that his payoff is a fixed price in case of no disclosure and increases in his type if he discloses. As a result, the incentive to disclose is increasing in type. The threshold equals the price the manager would obtain upon no disclosure, $P(\varnothing)$. Given that the price equals the market's expected value conditional on no disclosure, we have

$$P(\varnothing) = \frac{(1-\pi)E[V_1] + \pi * Pr(V_1 < P(\varnothing))E[V_1|V_1 < P(\varnothing)]}{(1-\pi) + \pi * Pr(V_1 < P(\varnothing))}.$$

One can express the equilibrium threshold, x^*, as a fixed point. Consider the expected value of V_1 conditional on no disclosure and the disclosure threshold being x:

$$\hat{P}(\varnothing, x) \equiv \frac{(1-\pi)E[V_1] + \pi * Pr(V_1 < x)E[V_1|V_1 < x]}{(1-\pi) + \pi * Pr(V_1 < x)}.$$

Then, x^* is a solution to $\hat{P}(\varnothing, x) = x$. This model was first introduced by Dye (1985). Jung and Kwon (1988) show the existence and uniqueness of such a fixed point (and hence of equilibrium). Acharya et al. (2011) show an alternative characterization named the "minimum principle" to which we will refer later:

$$P(\varnothing) = \min_x \hat{P}(\varnothing, x). \tag{1}$$

Note that an immediate implication of this condition is the uniqueness of the equilibrium (because if the minimum is also a fixed point, then $\hat{P}(\varnothing, x)$ has only one minimum; see Figure 1 for illustration). Moreover, for any x such that $x < \hat{P}(\varnothing, x)$, we have $x < x^* < \hat{P}(\varnothing, x)$, and for any x such that $x > \hat{P}(\varnothing, x)$, we have $\hat{P}(\varnothing, x) > x^*$: see Figure 1.

Finally, consider the expected payoff for the manager. With probability π he is informed, and his payoff is $\max\{P(\varnothing), V_1\}$, because he discloses if and only if doing so increases his price. With probability $1-\pi$, the manager has no discretion, and his payoff is $P(\varnothing)$. The law of iterated expectations implies that the manager's average payoff equals the firm's average value based on the distribution of F, that is,

$$E[V_1] = (1-\pi)P(\varnothing) + \pi \int_{V_1} \max\{P(\varnothing), V_1\} dF.$$

3.2 Existence and Uniqueness in the Dynamic Model

We now turn to the dynamic game and provide preliminary results about the existence, uniqueness, and structure of the equilibrium:

Lemma 1.

(i) In any equilibrium the agent follows a threshold strategy.

(ii) An equilibrium exists.

(iii) For π small enough the equilibrium is unique.

The term "threshold strategy" in (i) means that for any given history H_{t-1}, if a type $V_t = v$ discloses with positive probability, then all higher types $V_t > v$ find it strictly optimal to disclose. Thus, there exists a threshold x_t (that depends on the history, H_{t-1}) such that types above it disclose and types below it do not.

The proof of part (i) is as follows. Consider any equilibrium and any history H_{t-1}. If the agent discloses today $V_t = v$, his expected equilibrium payoff (starting today) is $v \sum_{s \geq t} w_s$ (because values are a martingale and the equilibrium prices are correct on average). Let $G(v)$ denote the payoff of the agent with $V_t = v$ if he does not disclose today, and subsequently follows an optimal disclosure strategy (this payoff includes the no-disclosure price today and the expected future prices given optimal disclosure in the future).

Clearly, $G(v)$ is increasing since a type $v' > v$ can follow the strategy of disclosing for the same value increments as those for which v discloses. If v' does so, he gets the same prices until the first disclosure and a strictly higher (by $v' - v$) price upon disclosure (and a strictly higher continuation payoff).

By a similar rationale $G(v') - G(v) < (v' - v) \sum_{s \geq t} w_s$ because the lower type can mimic the disclosure policy of the higher type (as before): in the event of no disclosure the lower type gets the same price as the higher type, and in the event of disclosure gets only $(v' - v)$ less. Since this is not the optimal strategy for the lower type, and no disclosure happens with positive probability on the equilibrium path (in particular, in period t), the strict inequality follows.[5]

[5] We assume that $w_s > 0$ for some $s \geq t$ to make the problem non-trivial. That assumption is sufficient for the strict inequality since there is a strictly positive probability that the manager will not be able to disclose in periods t to s. In that event, both types get the same payoff.

Rearranging this inequality we get that

$$G(v') - v' \sum_{s \geq t} w_s < G(v) - v \sum_{s \geq t} w_s. \tag{2}$$

So indeed, if the lower type v weakly prefers to disclose at t, which requires that $G(v) - v \sum_{s \geq t} w_s \leq 0$, the higher type strictly prefers to disclose.

We prove part (ii) (existence) in the appendix via relatively standard arguments that use Berge's theorem and a fixed point theorem. Roughly, think about every history of length $\{1, ..., T\}$ in which there was no disclosure. Pick some probabilities of disclosure in every period for each of these histories. They imply unique disclosure thresholds and, by Bayes' rule, unique no-disclosure prices that are continuous in the probabilities. Given those prices, find best-response disclosure thresholds for the agent. By Berge's theorem, the best-response thresholds are continuous in prices. This procedure creates a continuous mapping from probabilities of disclosure to probabilities of disclosure. That mapping has a fixed point, and that fixed point pins down the equilibrium (prices and thresholds).

Regarding uniqueness (part iii; also see the appendix), note that the analysis of equilibria in the multi-period game is more involved than in the static case we discussed above. In the static model, the uniqueness of the equilibrium can be shown without making any additional assumptions. In the dynamic model, uniqueness is not guaranteed. The main reason for this is that, unlike in the static case, in periods $t < T$ the expected no-disclosure payoff of type V_t depends on V_t. Even though the current-period no-disclosure price does not depend on it, the expected future prices do. In particular, considering two types that do not disclose today, a higher type today expects higher values in the future and hence higher expected future payoffs. Moreover, expected future prices depend on the market's beliefs about past disclosure thresholds, and this creates the possibility of multiple equilibria.

To see how the proof works when π is low, fix some history H_{t-1}. Let $\mathbf{x}_t$ denote the sequence of thresholds up to time t that the market believes the agent follows in his disclosure strategy at times $\{0, ..., t\}$. At t, the equilibrium threshold x type has to be indifferent between revealing $V_t = x$ and hiding it. As we argued before, if he reveals it, his expected payoff is $x \sum_{s \geq t} w_s$. If he does not disclose, the current period price is $P_t(H_{t-1}, \varnothing)$, and then there are continuation prices given equilibrium beliefs and optimal disclosure in the continuation game. Thus the equilibrium condition is

$$x \sum_{s=t}^{T} w_s = w_t \hat{P}_t(H_t, \mathbf{x}_t) + \sum_{s=t+1}^{T} w_s E[P_s | H_t, V_t = x]. \tag{3}$$

To claim uniqueness, we compare the derivative of the LHS of (3) with respect to x to that of the RHS. In order to rule out the possibility of more than one solution, it is sufficient to show that the derivative of the RHS is always higher. The derivative of the LHS is constant and equals $\sum_{s=t}^{T} w_s$. One might be tempted to conclude that our bound of the derivative of $G(v)$ in part (i) implies that the derivative of the RHS of (3) is lower than $\sum_{s=t}^{T} w_s$, which would imply the uniqueness of the equilibrium. This is not true because when we argue about $G(v)$, we take the continuation equilibrium as given. To find the equilibrium, we also need to account for how continuation prices change when we change the threshold at time t. Moreover, as x_t changes, future optimal thresholds change as well, further affecting prices. We show that when π is not too high, these additional effects get small as well.

To see the intuition for that claim, note that as $\pi \to 0$, the RHS of (3) converges to V_τ where τ is the last disclosed value (or $\tau = 0$ if no value was disclosed yet). Moreover, the derivative of the RHS converges to zero as $\pi \to 0$ because the probability that the market assigns to the agent knowing the value in any period between τ and T goes to zero. Hence, the sensitivity of future prices to any past thresholds goes to zero too. Hence, at least for small π, there is a single solution to that equilibrium condition.

In the rest of the paper, we assume that the parameters of the game are such that the equilibrium is unique. This simplifies the arguments. That said, similar results can be derived for the set of equilibria, in particular, if we focus on the equilibria that are the highest and the lowest solutions to (3) in any period.

4 Excess Disclosure

We now present our main result. We argue that in the equilibrium of the dynamic model, the manager reveals more information than he would in a one-period model.

In a one-period model, the manager reveals information if and only if disclosure yields him a higher price than the no-disclosure price. We show that in the dynamic model there is more disclosure. Namely, we show two results: first, the dynamic-model no-disclosure threshold is below the static-model disclosure threshold. Second, since in the static model the no-disclosure price satisfies the simple minimum principle, in the dynamic model, the manager discloses some information even though it decreases his current price to below the no-disclosure price. This equilibrium behavior might be seen as "excessive" disclosure if we tried to understand it through the lens of a static model.

The reason the manager may disclose some information that hurts his current price is that he cares about future prices, and we can prove that disclosing information today improves expected prices tomorrow. It may at first seem counterintuitive that additional disclosure of a type below the expected no-disclosure type could improve future prices: after all, the manager and the market are risk-neutral, and values and prices are martingales. Nevertheless, we show that in equilibrium, by disclosing V_t today, the manager reduces the uncertainty that the market will face in the future, and that improves his future no-disclosure prices. Thus it may seem that the manager is risk averse

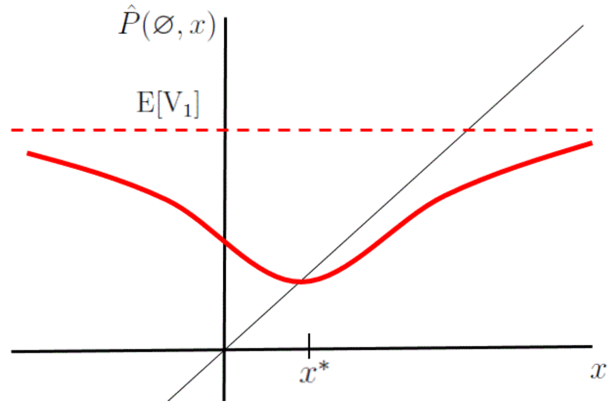

Figure 1: The no-disclosure price in a one-period model as a function of the threshold x. The unique equilibrium threshold is x^* and, as the graph shows, it equals the minimum no-disclosure price.

despite actually being risk-neutral. Given that disclosure is a call option, it may appear that the manager should prefer more uncertainty. However, the "strike price" of that option is the no-disclosure price that goes down with the market's uncertainty (because the market is more worried about the manager hiding negative information) and so the value of the manager's option increases with disclosure.

Our proof strategy is based on the following logic. We assume by way of contradiction that the manager follows a myopic disclosure strategy. Given these market beliefs, we then show that the threshold type would actually strictly prefer to disclose. This leads to a contradiction. A threshold that would make him indifferent has to be lower (the solution to (3) has to have a lower x than the one that solves $x_t = \hat{P}_t(H_t, \mathbf{x}_t)$, which defines the myopic threshold). This also implies that the equilibrium no-disclosure price in the dynamic model is higher than in the static model. That is implied by the minimum principle we discussed before: the myopic price is the lowest price over all possible disclosure policies in period t (see Figure 1).

In a two-period model, this result is relatively simple to prove. One can use the lemma of Jung and Kwon (1988).

Lemma 2. (Jung and Kwon 1988) *Consider a one-period model and let* x_F^* *denote the threshold when the prior distribution is* F*. If* F *dominates* G *in the second-order stochastic dominance sense, then* $x_F^* > x_G^*$.

With this lemma at hand, suppose that in period $t = 1$ (when $T = 2$) in equilibrium the manager follows a myopic strategy: disclose if and only if $P_1(V_1) = V_1 > P_1(\varnothing)$. To see why the threshold type would strictly prefer to disclose, consider the market beliefs in the second period. From the point of view of the threshold type there are two relevant histories: $H_1 = (x_1^{myopic}) = (P_1(\varnothing))$ or $H_1 = (\varnothing)$. That is, he either discloses or not.

If the manager followed a myopic threshold in the first period, the market beliefs about V_2 would have the same mean (x_1^{myopic}) after both histories. However, if he does not disclose in the first period, market beliefs would be more dispersed. In fact, the market beliefs upon no-disclosure are dominated in the second-order stochastic dominance sense by the market beliefs upon disclosure. Therefore, Lemma 2 implies a lower no-disclosure price in the second period if he does not disclose in the first period: $P_2(\varnothing, \varnothing) < P_2(x_1^{myopic}, \varnothing)$ (recall that in the static game thresholds are equal to no-disclosure prices and so the ranking of thresholds implies a ranking of prices). Since the distribution of values in the second period, V_2, is unaffected by the disclosure in the first period, we conclude that the expected payoff of the threshold type of the manager would be strictly higher in the second period if he disclosed in the first period, a contradiction. Instead, the first-period equilibrium threshold has to be lower than the myopic threshold (and the first-period no-disclosure price higher than under myopic disclosure because of the minimum principle).[6]

[6] This reasoning seems to directly establish only that the equilibrium threshold has to be lower than the equilibrium no-disclosure price. But inspecting Figure 1 (which follows from the minimum principle in Section 3.1), we see that the threshold is below the no-disclosure price if and only if $x_1 < x_1^{myopic}$.

For the general case of an arbitrary number of periods, we need a different proof strategy because disclosure at time t affects prices in all future periods, and except in period T, disclosure is not myopic.

Suppose that the current time is t and the history is H_{t-1}. The current-period equilibrium threshold is $x_t^*(H_{t-1})$. Define by $x_t^{myopic}(H_{t-1})$ an equilibrium threshold in a single-period game with the same prior probability distribution over V_t as in the equilibrium of the dynamic game.[7] Our main result is that the equilibrium disclosure thresholds are lower than the myopic thresholds and hence lower than the no-disclosure prices.

Theorem 1. *For every $t < T$ and every history H_{t-1},*

$$x_t^*(H_{t-1}) < x_t^{myopic}(H_{t-1}) < P_t(H_{t-1}, \varnothing).$$

Proof. Fix an arbitrary history H_{t-1} and the equilibrium thresholds before t. In this proof we abuse notation and let $P_t(\varnothing) \equiv P_t(H_{t-1}, \varnothing)$ to suppress the dependence on H_{t-1}. Let $h(v)$ denote the expected future payoff (not including time t) for a type $V_t = v$ conditional on him not disclosing at t.[8]

We first show that h is convex. We argue that for all $\Delta > 0$,

$$h(v) < \frac{1}{2}(h(v + \Delta) + h(v - \Delta)).$$

Consider the optimal continuation strategy of type $V_t = v$ as a function of future increments of value. Suppose that types $(v + \Delta)$ and $(v - \Delta)$ imitate it. Under the imitation strategies, the average payoff of the two types equals $h(v)$. The reason is that by disclosing when type v would disclose, types

[7] The prior depends on the history H_{t-1} and past equilibrium disclosure thresholds. As we discussed in the one-period model, given this prior, $x_t^{myopic}(H_{t-1})$ is equal to the minimum of no-disclosure prices over all period t disclosure policies.

[8] Again, this payoff depends on H_{t-1} but, for clarity, we suppress this dependence in the notation.

$(v+\Delta)$ and $(v-\Delta)$ get the same price on average. By not disclosing when type v does not disclose, types $(v+\Delta)$ and $(v-\Delta)$ get the same no-disclosure price. Since exactly mimicking type v is not the optimal strategy for either $(v+\Delta)$ or $(v-\Delta)$, the inequality follows.

Let $E[h(V_t)]$ be the average continuation payoff of all types that do not disclose in equilibrium at t. Based on the convexity of h, Jensen's inequality implies

$$h(P_t(\varnothing)) < E[h(V_t)|d_t = \varnothing], \tag{4}$$

where $P_t(\varnothing) = E[V_t|d_t = \varnothing]$ by rational expectations, i.e., because the equilibrium no-disclosure price is the average of types that do not disclose at t.

We can now prove the claim in the theorem. The equilibrium threshold x_t^* has to satisfy:

$$x_t^* \sum_{s \geq t} w_s = w_t P_t(\varnothing) + h(x_t^*). \tag{5}$$

Suppose by way of contradiction that there exists an equilibrium in which at $t < T$, $x_t^* \geq P_t(\varnothing)$. First, recall that (2) implies that if we compare the continuation payoffs of the threshold type and the type equal to the no-disclosure price, then, because the lower type can mimic the higher type's strategy, we get

$$h(x_t^*) - h(P_t(\varnothing)) \leq (x_t^* - P_t(\varnothing)) \sum_{s>t} w_s. \tag{6}$$

Second, Jensen's inequality (4) and the equilibrium conditions imply that

$$h(P_t(\varnothing)) < E[h(V_t)] = P_t(\varnothing) \sum_{s>t} w_s, \tag{7}$$

where the equality follows from the value process being a martingale: since the market believes that all the types that do not disclose have an average value $P_t(\varnothing)$, it follows that averaging future prices at any future date is also $P_t(\varnothing)$.

Combining (6) and (7) we get

$$h(x_t^*) < x_t^* \sum_{s>t} w_s,$$

which contradicts (5). Since we know that an equilibrium exists, it must therefore be that in equilibrium $x_t^* < P_t(\varnothing)$.

Finally, recall that the minimum principle implies that for any $x_t < \hat{P}_t(\varnothing, x_t)$ we have $x_t < x_t^{myopic} < \hat{P}_t(\varnothing, x_t)$ (see Figure 1), as claimed in the statement of the theorem.

□

4.1 The "Suspicious Belief Principle"

A key property in static disclosure models is the "minimum principle". According to this principle, whenever the agent does not disclose information, in equilibrium the market assumes the worst possible scenario. For instance, where it is commonly known that the agent is informed, as in Grossman (1981) and Milgrom (1981), the equilibrium price upon no-disclosure is the lowest value. Moreover, when it is not commonly known that the agent is informed, the no-disclosure price is the lowest expected value that any disclosure policy can support, see (1) and Figure 1.

In contrast, in our dynamic model, the disclosure policy does not minimize the first-period no-disclosure price (Theorem 1). But does the spirit of the "minimum principle" survive in any way in our game?

Notice first that no one belief minimizes all no-disclosure prices in our dynamic setting. As a result, if we wanted to find a market belief that minimizes the sum of no-disclosure prices weighted by w_t and the probability the agent reaches t without disclosing, the minimizer would depend on the disclosure strategy (via the probabilities).

As a result, instead of the "minimum principle," a more general property holds in our game. We call it the "suspicious belief principle." This new principle simplifies to the "minimum principle" if the agent assigns positive w_t to only one period.

To define the "suspicious belief principle," let $P_s^{\hat{\sigma}}(\varnothing)$ be the no-disclosure price at time s, given that there was no disclosure up to time s and the market believes the agents follows $\hat{\sigma}$. Next, suppose that the agent follows σ, but the market believes that he follows $\hat{\sigma}$. For these two strategies, define the weighted sum of expected no-disclosure prices as:

$$\phi(\sigma, \hat{\sigma}) = E\left[\sum_{s=1}^{\tau(\sigma)-1} w_s \cdot P_s^{\hat{\sigma}}(\varnothing)\right], \tag{8}$$

where $\tau(\sigma)$ is a random variable which is equal to the first disclosure period given σ and the expectation is with respect to that random variable.[9]

Note that $\hat{\sigma}$ affects $\phi(\sigma, \hat{\sigma})$ only via the no-disclosure prices and σ affects it only via the stopping times. As we discussed above, changing $\hat{\sigma}$ can increase the no-disclosure prices in some periods and reduce them in other periods. For example, the strategy that minimizes $P_1^{\hat{\sigma}}(\varnothing)$ (the static model strategy) does not minimize $P_2^{\hat{\sigma}}(\varnothing)$.

With this notation, $\phi(\sigma, \sigma)$ is the sum of expected no-disclosure prices if the market believes correctly that the agent follows σ. By contrast, $\phi(\sigma, \hat{\sigma})$ is the sum for the same agent strategy but when the market believes incorrectly that he follows $\hat{\sigma}$. The following proposition states that σ^* is an equilibrium strategy if and only if for *any* potential strategy σ, the equilibrium belief σ^* is *worse* (in terms of the sum of no-disclosure prices) than if the market believed σ.

[9] To simplify notation, we have defined $\phi(\sigma, \hat{\sigma})$ only from the perspective of $t = 1$. However, recall that our model "resets" after every disclosure: if the agent discloses $V_t = v$ then the continuation equilibrium is also an equilibrium of a game with horizon $T - t$ and starting value $V_0 = v$. For periods following disclosure the function in (8) is re-computed by replacing $s = 1$ with $s = t$.

Proposition 1. *Consider the beginning of the game or any history following disclosure. σ^* is an equilibrium strategy if and only if for any strategy σ,*

$$\phi(\sigma, \sigma^*) \leq \phi(\sigma, \sigma).$$

Our proof shows that if this condition is violated for some σ, then this σ would be a profitable deviation. The intuition starts with noting that if the market observed the deviation to σ, the average price in every period would be v_0, as in equilibrium (since the market updates correctly). However, since the market does not observe deviations, if $\phi(\sigma, \sigma^*) > \phi(\sigma, \sigma)$, the expected deviation payoff would be strictly higher.

One special case of our dynamic setting is when the agent cares only about the last period price. In this case there is one belief that minimizes the value upon no disclosure for all strategies as in the static model, and therefore the minimizing strategy is the unique equilibrium strategy. In other words, in the special case where the agent cares only about the final price, the condition of Proposition 1 implies the familiar "minimum principle":

Corollary 1. *If $\forall s \leq T-1,\ w_s = 0$, and $w_T = 1$, then*

$$\sigma^* = \arg\min_{\sigma} P_T^{\sigma}(\varnothing). \tag{9}$$

This corollary has an interesting implication regarding the information revealed about the final value V_T. Consider an arbitrary weight profile $\{w_t\}$ and compare it to the case in which the agent cares only about the final price. The above corollary implies that more information is being revealed regarding the final value in the latter case. The agent discloses more values when he cares only about the final price. As we shall see in section 5.1, this is also relevant for the hypothetical case in which the agent can disclose information in a non-timely manner.

5 Variants

Our goal in this section is to examine alternative assumptions to understand better what drives our results. We first consider a variant in which the manager can disclose stale information. We then consider a variant in which information arrives in the form of a Poisson process of information events. The agent's ability to disclose coincides with the realization of these events.

5.1 Allowing for Stale Disclosure

A key feature of our model is that the manager can disclose only timely information. For an intuitive understanding of how this restriction affects equilibrium information disclosure, we now examine a variant of the model where the agent can disclose both current and old (stale) information. There are two interpretations of this variant of the model: either the information becomes long-lived in the sense that even with delay it can be credibly revealed, or the requirement to disclose information in a timely manner is removed.

We analyze a two-period model and show that lifting the restriction for timely disclosure (or, in the opposite direction, requiring timely disclosure), has an ambiguous effect on the amount of information disclosed in equilibrium. On the one hand, allowing the agent to disclose information with delay reduces the likelihood of early disclosure of V_1 at $t = 1$. On the other, this is offset by more disclosure at a later date: we show that the threshold of disclosure of the first signal goes down over time and in the second period is below the threshold in our original model.

We start with a comparative static result for our main model (when the manager cannot disclose stale information). We examine the effect of changing the weight on the first period, w_1. We shall maintain the assumption that π is sufficiently small such that the equilibrium is unique.[10] We argue

[10]A similar result can be obtained for the case of multiple equilibria, but comparative

that:

Lemma 3. *In the two-period model, when the disclosure of stale information is not allowed, the threshold for disclosure in the first period is increasing in* w_1.

While we provide the formal proof in the Appendix, the economic intuition is clear. The larger the weight w_1, the less important the continuation payoff, and hence the equilibrium disclosure strategy gets closer to being myopic.

Suppose now that the manager can disclose stale information. Specifically, in addition to the disclosure opportunities in our main model, he can disclose V_1 at $t = 2$. This disclosure is relevant only when he does not also disclose V_2.

We first note that this possibility leads to multiple equilibria in which off-equilibrium beliefs play a major role. In particular, there exists an equilibrium in which stale information is never disclosed. It is based on negative off-equilibrium beliefs (for example, if the agent reveals V_1 in period 2, the market believes that V_2 has the smallest possible realization given V_1). However, some natural refinements eliminate this equilibrium. In particular, consider a variant of a "trembling-hand" that we define as follows.

Definition 2. *An* $\epsilon -$ *game is a game in which at each point in time, any piece of evidence is leaked with probability* ϵ*. The probability of this event is independent of (i) the value of this and other pieces of evidence and (ii) leakage and strategic release of other pieces of information.*

Definition 3. *A trembling-hand equilibrium is the limit of equilibria of* $\epsilon -$ *games when* ϵ *goes to zero.*

The trembling we consider implies that conditional on V_1 being leaked but not V_2, the manager discloses V_2 strategically. An alternative refinement

statics become clearer when the equilibrium is unique.

would be that with some small probability the manager is truthful. Thus, when only V_1 is leaked, it must be that the manager does not know V_2, and V_1 represents the entire truth. This is considered by Hart et al. (2017) who name this "truth leaning". One can show that the equilibrium outcome is the same under both refinements.

When the manager can disclose V_1 at $t = 2$ there are time $t = 2$ disclosure thresholds for both V_1 and V_2. We denote by $x^{*,stale}_{1,1}, x^{*,stale}_{1,2}$ the equilibrium thresholds for disclosure of V_1 at $t = 1$ and at $t = 2$, respectively. Similarly, we denote by $x^{*,stale}_{2,2}$ the equilibrium threshold for disclosure of V_2 at $t = 2$ when the agent has not disclosed V_1 and $x^{*,stale}_{2,2}(v)$ when he has disclosed $V_1 = v$. The equilibrium is, as before, a collection of these thresholds and history-dependent prices such that the thresholds are optimal given prices, prices are consistent with Bayes' rule and the equilibrium strategy on and off the path, and prices are consistent with the trembles for off-path disclosures.

We argue that in a trembling-hand equilibrium, stale information is treated in the same way as information that is revealed on time. That is, at $t = 2$ the market price when only V_1 is revealed is independent of whether it is revealed at $t = 1$ or $t = 2$. The manager's disclosure strategy is myopic. We show in the Appendix that:

Lemma 4. *In the two-period model with stale information, there exists a unique trembling-hand equilibrium. In this equilibrium, the disclosure strategy is myopic. The manager releases information at time t if and only if it leads to a higher price at time t than the no-disclosure price. Moreover, the disclosure of V_1 has the same effect on the price at $t = 2$ regardless of whether it was disclosed at $t = 1$ or $t = 2$.*

5.1.1 Comparison of Information Disclosure

We consider the disclosure of V_1. We argue that there is less disclosure in the first period but overall more disclosure when it can be disclosed at $t = 2$.

That is, the stale disclosure of V_1 at $t = 2$ more than compensates for the lower level of disclosure at $t = 1$.

Proposition 2. *Suppose that $w_1, w_2 > 0$. Then, $x_{1,2}^{*,stale} < x_1^* < x_{1,1}^{*,stale}$.*

Proof. The proof of why $x_1^* < x_{1,1}^{*,stale}$ follows from Lemma 4, which implies $x_{1,1}^{*,stale} = x_1^{myopic}$, and Theorem 1 which implies that $x_1^* < x_1^{myopic}$.

The proof of why $x_{1,2}^{*,stale} < x_1^*$ follows from a few observations. First, we argue that in the stale model, the decision to disclose V_1 can be made independently of V_2. The reason is that disclosure of V_1 is relevant to the price only in the case where V_2 is not disclosed. The manager's disclosure strategy can be viewed as disclosing V_1 if it leads to a higher price compared to not disclosing anything, assuming that V_2 will not be disclosed. Then, if he learns that V_2 is even higher than the price after deciding about V_1, he discloses it. Second, $x_{1,2}^{*,stale}$ equals x_1^* when $w_1 = 0$. That is, it equals the time $t = 1$ disclosure threshold when stale information is not allowed and the manager cares only about the time $t = 2$ price. This follows from the first observation and Lemma 4 (that the disclosure of V_1 has the same effect on the price at $t = 2$ regardless of whether it was disclosed at $t = 1$ or $t = 2$). Finally, the claim that $x_{1,2}^{*,stale} < x_1^*$ follows immediately from the second observation and Lemma 3 (that x_1^* increases in w_1). □

It follows from the proposition that there is a range of (low) values of V_1 that are disclosed only if stale information is allowed, and not otherwise.[11] A similar result holds with regard to V_2:

Proposition 3.

$$x_2^{*,stale} < x_2^*.$$

Proof. In the proof of Proposition 2 we used the idea that when stale information is allowed, the agent disclosure policy in the second period is just

[11] Nevertheless, this disclosure is not timely since an agent discloses such a value only in the second period (when it is a stale information).

like the disclosure policy of an agent in our main model where $w_1 = 0$ and $w_T = 1$. Therefore, the proposition follows directly from Corollary 1, Lemma 3, and the observation that in the last period in both models the agent is myopic. To see this, note that in the stale model, if there is no disclosure in any period, the last period is the minimum price over all possible disclosure strategies and it is the same price as in the main model if $w_1 = 0$. However, if $w_1 > 0$ then the first-period strategy is different (Lemma 3), and hence the last period price is higher. Since the last period disclosure threshold is equal to the no-disclosure price, it must also be higher. □

To summarize, there are low values of V_1 and low values of V_2 that are never disclosed if stale information is not allowed, but might be disclosed if stale information is allowed. In this sense, more information is disclosed when stale information is allowed to be disclosed. Nevertheless, as Proposition 2 shows, low values in the first period are disclosed only as stale information, i.e., at period 2.

A final remark applies to the case of more than two periods. Solving the game with stale information for longer horizons is quite difficult. The main difficulty is in proving that the equilibrium strategy is myopic (as we have done in the two-period model). Nevertheless, we have verified in many examples using numerical solutions that there exists a myopic equilibrium (i.e., the agent discloses information to maximize the current price.). When the equilibrium disclosure is myopic, our main conclusion holds even with longer horizons. Namely, when the disclosure of stale information is allowed, there is less early disclosure. Yet, over time, disclosure thresholds go down, and eventually, more information is revealed.

5.2 Disclosure When Value Changes

Consider now a variant of the model in which the value changes in some periods only, and when it does, the manager has verifiable information. In particular, in every period, with probability π, an information event occurs. If so, the value, V_t, increases by ΔV_t, and the manager has information he can disclose at that time (as before $\{\Delta V_t\}$ are i.i.d. random variables). With probability $1-\pi$, the value does not change so that $V_t = V_{t-1}$ and the manager cannot prove that the value has not changed. One can think about two different cases for such a game. In the first, the manager can disclose the cumulative value V_t. In the other, he can only disclose the current increment ΔV_t.

When the manager discloses V_t, note that the model is not immediately a special case of our main model. We have assumed so far that the ability to disclose is independent of the value. Nevertheless, one can verify that our results from Section 4 continue to hold with essentially the same proofs, because this model maintains that the probability the agent can disclose in the future is independent of the current value. The proof for why there is excess disclosure is also valid for this case. The agent follows a threshold strategy where the thresholds are lower than the myopic one.

When instead the manager discloses the increments ΔV_t, then the agent becomes myopic. In period t the price equals the time $t-1$ plus the expected value of the increment conditional on the time t disclosure, $P_t = P_{t-1} + E(\Delta V_t|d_t)$. It follows that P_{t-1} does not affect the disclosure decision. Hence, the disclosure decision is the same as the static model discussed above, where the agent decides whether to disclose ΔV_t.

This comparison shows that our model yields different results depending on the nature of the information the agent can disclose. If it is just current increments, disclosure does not depend on past decisions. If it is cumulative value, previous disclosure decisions do affect current disclosure. It matters whether the market knew V_{t-1} so $P_{t-1} = V_{t-1}$ or P_{t-1} was based on partial information (even if the price last period was the same).

6 Conclusions

We have studied a dynamic model of voluntary information disclosure in a setup where the firm's true value follows a random walk. The manager occasionally learns verifiable information and chooses what to disclose to maximize a weighted sum of stock prices.

We have derived two main results. First, the equilibrium disclosure threshold is strictly lower than the no-disclosure price. The manager in equilibrium sometimes discloses information even though doing so reduces the current stock price. We refer to this result as "excess disclosure" and it runs counter to most existing models. The intuition is that the manager has an incentive to reduce the future market's uncertainty (despite all agents being risk-neutral). Second, we have shown to what extent the "minimum principle" from static models extends to our dynamic model. The "minimum principle" states that when the market sets the price upon no disclosure, it assigns the most pessimistic beliefs about the agent's strategy. This property does not hold in the dynamic model and is replaced by a more general "suspicious belief principle" we discovered.

Finally, focusing on two-period models, we compare the timely disclosure model to a stale information model (i.e., where the manager can disclose information with delay). We show that ranking the amount of information disclosed in the two models is complex. Specifically, in the stale model, where the agent is not required to disclose information in a timely manner, there is less disclosure in the first period but more disclosure overall. An interpretation of this result is that regulation in response to managers delaying information disclosure can unintentionally result in less rather than more information reaching the market.

References

Acharya, V. V., P. DeMarzo, and I. Kremer (2011). Endogenous information flows and the clustering of announcements. *American Economic Review 101*(7), 2955–79.

Aghamolla, C. and B.-J. An (2019). Voluntary disclosure with evolving news. *Available at SSRN*.

Bertomeu, J., A. Beyer, and R. A. Dye (2011). Capital structure, cost of capital, and voluntary disclosures. *The Accounting Review 86*(3), 857–886.

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

Einhorn, E. and A. Ziv (2008). Intertemporal dynamics of corporate voluntary disclosures. *Journal of Accounting Research 46*(3), 567–589.

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

Grossman, S. J. and O. D. Hart (1980). Takeover bids, the free-rider problem, and the theory of the corporation. *The Bell Journal of Economics*, 42–64.

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

Hart, S., I. Kremer, and M. Perry (2017). Evidence games: Truth and commitment. *American Economic Review 107*(3), 690–713.

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

Kreps, D. (2020). *Microeconomic Foundations, II: Imperfect Competition, Information, and Strategic Interaction, Appendix 12*. draft.

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

Pae, S. (2005). Selective disclosures in the presence of uncertainty about information endowment. *Journal of Accounting and Economics 39*(3), 383–409.

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

Shin, S. H. (2006). Disclosure risk and price drift. *Journal of Accounting Research 44*(2), 351–379.

Appendix

Proof of Lemma 1: We have proven part (i) in the text.

Part (ii) (existence): Start with an auxiliary game: for each $\tau \in \{0, ..., T-1\}$ define the following game. First, $V_\tau = 0$ and the game starts at $t = \tau + 1$. Second, in this game, V_t has the increments distributed in the same way as in the original game. The agent observes each V_t with the same probability as in the original game and decides whether to disclose in every period. If he does not disclose at t, he obtains a payoff flow $w_t P_s(H_t^{\varnothing})$ where $H_t^{\varnothing}$ is the history of no disclosure between $\tau + 1$ and t. If he discloses at t the game ends and he gets a final payoff of $V_t \sum_{s \geq t} w_s$. An equilibrium of the auxiliary game for a fixed τ is a sequence of disclosure thresholds x_t^* and no-disclosure prices P_t for all $t \in \{\tau + 1, ...T\}$, such that the thresholds are optimal given the prices and the prices are consistent with the thresholds and Bayes' rule.

If we can find the equilibria of this auxiliary game for every τ, then we can construct an equilibrium for the original game. To see this, note that if in the original game $V_\tau = v$ is disclosed, and if we make the continuation prices (until the next disclosure) and disclosure thresholds equal to the prices and thresholds from the auxiliary game plus v, then all incentives continue to be satisfied, and prices continue to be consistent with equilibrium strategies.

Now, fix any τ and consider an arbitrary vector of probabilities of disclosure in every period $t \in \{\tau+1, ...T\}$ conditional on not disclosing before (and conditional on having verifiable information in that period). These probabilities pin down uniquely the disclosure thresholds in all periods. Second, compute for that vector of disclosure probabilities the implied thresholds, and then the implied no-disclosure prices. These prices are continuous in the vector of probabilities. Third, for any arbitrary vector of prices $\tilde{P}_s$ for $s > \tau$, consider the best-response problem of the agent in the auxiliary game given those prices. The objective function of the agent is continuous in the prices and the probabilities. It follows from Berge's theorem that the best-response correspondence (from prices to optimal probabilities of disclosure) is upper semicontinuous. Moreover, we claim (below) that the best response is unique, and so the best response is a continuous function from the vector of prices to the vector of probabilities disclosure (implied by the optimal deterministic thresholds).

Putting these operations together (from the vector of probabilities of disclosure to no-disclosure prices using Bayes' rule, and then from no-disclosure prices back to probabilities of disclosure using the best response), we define a continuous mapping from conjectured probabilities of disclosure back to the vector of probabilities of disclosure. By Brouwer's fixed point theorem, this mapping has a fixed point. That fixed point is an equilibrium of the auxiliary game.

To see that the best-response disclosure policy is unique for any vector of no-disclosure prices $\tilde{P}_s$, note the following. First, the disclosure threshold at T is equal to $\tilde{P}_T$. Second, for any other period t, the disclosure threshold is independent of disclosure thresholds in previous periods and can be found by solving:

$$x_t \sum_{s\geq t} w_s = w_t \tilde{P}_t + h(x_t|\tilde{P}), \tag{10}$$

where $h(v|\tilde{P})$ is the expected optimal continuation payoff of type v if he does not disclose today, given the future no-disclosure prices.

The derivative of the LHS of (10) is $\sum_{s\geq t} w_s$ while the derivative of the RHS is strictly less. By the envelope theorem, the latter is equal to $\sum_{s>t}(w_s \Pr(\text{disclosure at } s \text{ or before } |V_t = x_t))$. Hence, either for all x_t the LHS is larger (so the best response is to disclose with probability 1) or the RHS is larger (so the best response is to disclose with probability 0) or there is a unique interior x_t that satisfies (10). The intuition for why the derivative of the RHS is smaller than $\sum_{s\geq t} w_s$ is analogous to the argument we made in the proof of part (i). Fix the optimal continuation strategy of type v' and consider a lower type $v < v'$. That type can mimic v' by disclosing for the same value increments. Conditional on disclosure their payoffs differ by $v' - v$ (weighted by $w's$) and conditional on no disclosure their payoffs are the same.

Part (iii) (uniqueness for small π): We argue that at least for small π, for an arbitrary history H_t, and for an arbitrary threshold strategy up to time t, $(x_1, ...x_t)$, there exists a unique threshold equilibrium for the sub-game for periods $s > t$. Let $\mathbf{x_t} \equiv (x_1, ..., x_t)$. In the proof we shall rely on the following observation. For all histories, thresholds, and $\tau \leq t$

$$\frac{\partial \hat{P}_t(H_t, \mathbf{x_t})}{\partial x_\tau}$$

converges uniformly (across all $\mathbf{x_t}$) to zero as $\pi \to 0$. The intuition is that x_τ is relevant only when there is no disclosure in periods $\tau, ... , t$. For small π, the market's equilibrium beliefs are that it is most likely that the manager could not disclose at τ. Thus, as π gets small, x_τ has a diminishing impact on no-disclosure prices.[12]

We prove the claim by backward induction on t, showing that for every history H_t, all arbitrary thresholds $\mathbf{x_t}$, and all periods $s > t$:

(a) There exists a unique continuation equilibrium with disclosure thresholds $x_s^*(H_s|\mathbf{x_t})$ and prices $P_s(H_s|\mathbf{x_t})$ for every history H_s consistent with H_t.

(b) $\forall \tau \leq t : \frac{\partial x_s^*(H_s|\mathbf{x_t})}{\partial x_\tau}$ converges uniformly (across all $\mathbf{x_t}$) to zero as $\pi \to 0$.

Before continuing with the formal proof, we should note that the intuition for (b) is similar to the intuition described above for why the no-disclosure prices become insensitive to thresholds as π gets small: optimal thresholds today depend on the current and future no-disclosure prices; if those prices are not sensitive to past thresholds, today's threshold also becomes insensitive to them.

We begin the proof by induction by letting $t = T - 1$. Part (a) of the hypothesis follows directly from our discussion of the one-period model above: in a single-period game, the equilibrium is unique for all prior distributions. Part (b) holds because, letting $H_T = (H_{T-1}, \varnothing)$, the threshold $x_T^*(H_{T-1}|\mathbf{x_{T-1}})$ (for disclosure of V_T) is a solution to

$$x_T = \hat{P}_T(H_T, \mathbf{x_T}),$$

[12]The claim can be proven by first expressing $\hat{P}_t(H_t, \mathbf{x_t})$ using Bayes' rule as a sum of terms that correspond to different combinations of signals the agent has observed so far. Each of these terms has a bounded derivative (uniformly for all $\mathbf{x_t}$ since $f_s(v)$ and $|vf_s(v)|$ are being uniformly bounded). Moreover, each such derivative is multiplied by π (other than the first term, which is that the agent received no signals, but the derivative of that term is zero), and so the limit derivative uniformly converges to zero.

so that by the implicit function theorem we have that

$$\frac{\partial x_T^*(H_{T-1}|\mathbf{x_{T-1}})}{\partial x_\tau} = \frac{\frac{\partial \hat{P}_T(H_T,\mathbf{x_T})}{\partial x_\tau}}{1 - \frac{\partial \hat{P}_T(H_T,\mathbf{x_T})}{\partial x_T}}.$$

The claim for T follows because the derivatives of prices uniformly converge to zero.

Now, consider an arbitrary t and an arbitrary history H_{t-1}. Let $H_t = (H_{t-1}, \varnothing)$. Suppose that time τ is the last time in which the manager has disclosed and let $d_\tau = V_\tau$ (if there was no disclosure, $\tau = 0$). Recall that the assumption that the increments ΔV are independent implies that without loss of generality we can assume that $V_\tau = 0$.[13] At time t the equilibrium condition is

$$x_t \sum_{s=t}^{T} w_s = w_t \hat{P}_t\left(H_t, \mathbf{x}_t\right) + \sum_{s=t+1}^{T} w_s E[P_s|H_t, V_t = x]. \tag{11}$$

The expression $E[P_s|H_t, V_t = x]$ in (11) is the expectation of prices at time s given the history H_t (hence, no disclosure in period t), past disclosure thresholds $\mathbf{x}_t$, and sequentially optimal disclosures in periods $(t+1, ..., T)$. That expectation depends on the past conjectured cutoffs $\mathbf{x}_t$ as well as the optimal cutoffs $\hat{x}_s^*(H_{s-1}, \mathbf{x_t})$ from the induction hypothesis.[14]

For uniqueness, note that the derivative of the LHS of (11) with respect to V_t is constant at $\sum_{s=t}^{T} w_s$, while as $\pi \to 0$, the derivative of the RHS goes to zero. The reason the derivative of RHS goes to zero is that the probability that the agent will be able to disclose in future periods goes to zero, and

[13] If we find an equilibrium when $V_\tau = 0$ we can then add to all prices and thresholds any constant v to get an equilibrium when $V_\tau = v$ and vice versa.

[14] Existence of the equilibrium for small π can be shown even more directly than in our general proof above, by noticing that for all thresholds, the (RHS) of (11) converges to $V_\tau \sum_{s=t}^{T} w_s$, where V_τ is the last disclosed value in H_{t-1}, and that V_τ is in the interior of the support of V_t for any $\tau < t$. Hence by the intermediate function theorem there exists a solution to (11).

we have assumed by the induction hypothesis that the derivatives of future thresholds go to zero as well. Since the derivatives of all future no-disclosure prices with respect to conjectured thresholds converge to zero as well, the derivative of the RHS indeed converges to zero. That establishes that for small π, (11) has a unique solution $x_t^*(H_{t-1}|\mathbf{x_{t-1}})$.

Finally, using (11) and the implicit function theorem we can show that the derivative of the period t equilibrium threshold with respect to any conjectured threshold at $\tau < t$ converges to zero as well:

$$\lim_{\pi \to 0} \frac{\partial x_t^*(H_{t-1}|\mathbf{x_{t-1}})}{\partial x_\tau} = \lim_{\pi \to 0} \frac{\sum_{s=t}^{T} w_s \frac{d\hat{P}_s(H_s, \mathbf{x_s})}{dx_\tau}}{\sum_{s=t}^{T} w_s (1 - \frac{\partial \hat{P}_s(H_s, \mathbf{x_s}))}{\partial x_t})} = 0,$$

where

$$\frac{d\hat{P}_s(H_s, \mathbf{x_s})}{dx_\tau} = \frac{\partial \hat{P}_s(H_s, \mathbf{x_s})}{\partial x_\tau} + \sum_{z=t+1}^{s} \frac{\partial \hat{P}_s(H_s, \mathbf{x_s})}{\partial x_z} \frac{\partial x_z^*(H_{z-1}|\mathbf{x_t})}{\partial x_\tau}$$

takes into account that past thresholds affect the RHS of (11) directly by changes in the believed distribution at the beginning of H_t and indirectly by changes in the future equilibrium thresholds.

The final observation to make is that in performing the induction step we need to move from functions $x_s^*(H_s|\mathbf{x_t})$ to functions $x_s^*(H_s|\mathbf{x_{t-1}})$ by substituting the unique equilibrium $x_t^*(H_{t-1}|\mathbf{x_{t-1}})$ in place of the arbitrary x_t (and adding $x_t^*(H_{t-1}|\mathbf{x_{t-1}})$ to the collection of the unique continuation thresholds now starting at time t). After that substitution, all these functions inherit the property (b) in the induction hypothesis. □

Proof of Proposition 1. First, recall that for σ^* to be an equilibrium strategy, it has to be that after disclosure of $V_t = v$, the continuation strategy is also an equilibrium in the game that starts with value $V_0 = v$ and has horizon $T - t$. So, without loss of generality we prove the statements for arbitrary T but only arguing about what σ implies about the first disclosure time, $\tau(\sigma)$.

Necessity: Suppose σ^* is an equilibrium strategy.

For any strategy σ, if the agent follows it and the market believes that he does, the expected payoff is the same because prices satisfy Bayes' rule:

$$\forall\sigma, \quad V_0 \sum_{s=1}^{T} w_s = \phi(\sigma, \sigma) + E\Bigg[V_{\tau(\sigma)} \cdot \sum_{s=\tau(\sigma)}^{T} w_s\Bigg], \tag{12}$$

where $V_{\tau(\sigma)}$ is a random value that is disclosed at the random time $\tau(\sigma)$.

Since this is true for every strategy, it is true also for the equilibrium strategy:

$$V_0 \sum_{s=1}^{T} w_s = \phi(\sigma^*, \sigma^*) + E\Bigg[V_{\tau(\sigma^*)} \cdot \sum_{s=\tau(\sigma^*)}^{T} w_s\Bigg]. \tag{13}$$

Now consider a deviation to some strategy σ_1 until the first disclosure and then following σ^*. For σ^* to be an equilibrium, this deviation cannot be profitable. This deviation yields the expected payoff:

$$\phi(\sigma_1, \sigma^*) + E\Bigg[V_{\tau(\sigma_1)} \cdot \sum_{s=\tau(\sigma_1)}^{T} w_s\Bigg]. \tag{14}$$

Hence, for the deviation not to be profitable we must have that (14) is weakly smaller than the right hand side of (12) (which has the same payoff as the equilibrium payoff in (13)). This implies $\phi(\sigma_1, \sigma^*) \leq \phi(\sigma_1, \sigma_1)$, as claimed.

Sufficiency: Suppose that for some σ^* the condition holds at time $t = 1$ and at every history following disclosure (where $\phi(\sigma, \hat{\sigma})$ is redefined to sum over times following the disclosure). Suppose by contradiction that σ^* is not an equilibrium so there exists a profitable deviation $\hat{\sigma}$. Moreover, there must exist at least one period such that it is profitable to deviate from σ^* to

$\hat{\sigma}$ up to $\tau(\hat{\sigma}) - 1$ (time of the first disclosure given the deviation strategy) and after that playing σ^*. Otherwise, by induction, deviating to $\hat{\sigma}$ would not be profitable. Without loss of generality, suppose that this deviation is profitable at $t = 1$.

For $\hat{\sigma}$ to be profitable, we must have:

$$V_0 \sum_{s=1}^{T} w_s < \phi(\hat{\sigma}, \sigma^*) + E\left[V_{\tau(\hat{\sigma})} \cdot \sum_{s=\tau(\hat{\sigma})}^{T} w_s\right]. \tag{15}$$

Applying (12) to $\hat{\sigma}$ we also get:

$$V_0 \sum_{s=1}^{T} w_s = \phi(\hat{\sigma}, \hat{\sigma}) + E\left[V_{\tau(\hat{\sigma})} \cdot \sum_{s=\tau(\hat{\sigma})}^{T} w_s\right]. \tag{16}$$

Putting these conditions together yields $\phi(\hat{\sigma}, \sigma^*) > \phi(\hat{\sigma}, \hat{\sigma})$, a contradiction. $\square$

Proof of Lemma 3: Suppose that $x_1^*(w_1)$ is the threshold for disclosure in the first period given weight w_1. It is based on the following indifference condition:

$$w_1 * (\hat{P}_1(\varnothing, x_1^*) - x_1^*) = w_2 * (x_1^* - E[\hat{P}_2(\varnothing, d_2, x_1^*, x_2^*)|V_1 = x_1^*]).$$

The LHS is the $t = 1$ gain from no-disclosure by the threshold type; we have shown in Section 4 that it is positive. The RHS is the time $t = 2$ expected loss from no-disclosure at $t = 1$ (the expectation is with respect to the optimal disclosure d_2). Now consider $w_1^2 < w_1^1$. If we keep the threshold for disclosure at $x_1^*(w_1^1)$ as we decrease w_1 from w_1^1 to w_1^2, the LHS becomes smaller, implying that the agent would strictly prefer to disclose. Therefore the equilibrium threshold in period 1 has to change with w_1. If instead, we take the threshold x_1 to be very low, then the LHS will continue to be positive, and the RHS will become negative. The intermediate value theorem (and our assumption that the equilibrium is unique) implies that $x_1^*(w_1^2) < x_1^*(w_1^1)$. □

Note that similar reasoning shows that the lemma is true for any $T \geq 2$.

Proof of Lemma 4: Consider the case in which the manager discloses V_1 in the first period, that is, $d_1 = \{V_1 = v\}$ for some value v. The second period is equivalent to a one-period model with one signal and with an initial value $V_0 = v$. The equilibrium in this sub-game exists and is unique. This also holds off-equilibrium when disclosing $d_1 = \{V_1 = v\}$ is not on the equilibrium path. Second, suppose that the manager reveals V_1 in period 2: $d_2 = \{V_1 = v\}$. Again, regardless of whether $d_2 = \{V_1 = v\}$ is on or off the equilibrium path, prices and the optimal disclosure strategy for V_2 are the same as in a one-period model with initial value $V_0 = v$. The claim for off-path disclosures of V_1 follows from our trembles that pin down the belief about the distribution of V_2 for any V_1 disclosed off-path. Hence, second-period prices when only V_1 is disclosed are the same regardless of when V_1 was disclosed. Prices are equal to the no-disclosure price in a one-period, one-signal model with initial value $V_0 = v$; we denote it by $P_2(V_1)$.

When there is no disclosure of V_1, the agent discloses V_2 if and only if it exceeds the no-disclosure price. Hence, given the disclosure (or the lack thereof) of V_1, the equilibrium behavior for the disclosure of V_2 is unique.

It remains to show that the disclosure thresholds (in periods 1 and 2) for V_1 are unique. First, notice that in every equilibrium the first-period threshold cannot be lower than the first-period no-disclosure price, as otherwise the manager would be better off withholding the threshold value and disclosing it only in the second period. On the other hand, the threshold cannot be greater than the first-period no-disclosure price because then the manager would be better off disclosing values in between those two values in the first period rather than hiding them. This implies that the first-period threshold is unique since the fixed point $x = \hat{P}_1(\varnothing, x)$ is unique. It also implies that this is the myopic threshold.

We now consider the disclosure of V_1 in the second period. We shall argue that there exists a unique no-disclosure price $P_2(\varnothing, \varnothing)$, which implies a unique disclosure threshold. Let $g(P)$ denote the "market profit" in the second period as a function of the no-disclosure price (that is, $g(P)$ is the difference between the expected type and the average prices paid given the no-disclosure price). Note that $g(P)$ is equal to the profit/loss for all types who do not disclose given P; any type who discloses is being correctly priced and leads to zero profit. As a result, $g(P) = 0$ if and only if P is equal to the expected value conditional on no disclosure. We argue that there exists a unique P^* for which $g(P^*) = 0$. That claim follows from three properties of g (and the intermediate value theorem). First, g is continuous. Second, $g(P) > 0$ for sufficiently low P and $g(P) < 0$ for sufficiently high P. Third, $g(P)$ decreases in P.

That $g(P) > 0$ for sufficiently low P follows since for very low P, the agent fails to disclose only when he does not receive a signal. The expected non-disclosing type is then the unconditional mean type, which results in underpricing. By contrast, when P is sufficiently high (for example, higher than the highest type) no types disclose. The value of these types is again the unconditional mean, and such high P results in overpricing. To show that $g(P)$ is decreasing, note that when we increase P from p_1 to $p_2 > p_1$ there are two effects to consider. First, we increase the payment to those who do not disclose at p_1 and continue not to disclose at p_2; this is a loss. Second, there are types who disclose when the price is p_1 and choose not to do so when the price is p_2. This again is a loss, as these types were priced correctly when they disclosed, but now they prefer to hide.

□