Dynamic Information Regimes in Financial Markets

Paul Glasserman Harry Mamaysky Yiwen Shen*

This version: July 2019

Abstract

We develop a dynamic, rational expectations model of information and asset prices in which investor information choices influence the level of available public and private information about fundamentals. We posit that as more investors become informed, more information about fundamentals becomes available. Two regimes emerge, one with high prices and low volatility, and one with low prices and high volatility. Information dynamics move the market between regimes, creating large market drops and rallies, with no change in fundamentals, but large changes in discount rates, reflecting changes in information asymmetry. We also study alternative dynamics in which an increase in the number of informed investors leads to greater public disclosure; this mechanism has a stabilizing effect. When calibrated to market data, the first mechanism suggests a role for information dynamics in financial crises; the second mechanism helps explain empirical findings on the market reaction to the loss of analyst coverage.

1 Introduction

Most research linking investor information acquisition and asset prices assumes a constant information environment. But why should the level of potentially available information remain constant in a market that is perpetually in flux? Changes in technology and regulation can generate persistent shocks to what an investor can learn about company fundamentals; and changes in what can be learned should influence investors as they decide whether to acquire costly information. Pushing this idea a step further, we investigate what happens when the information environment itself changes in response to investor demand for information. In other words, we posit that the news media, financial intermediaries, company executives, regulators,

*Glasserman: Columbia Business School, pg20@columbia.edu. Mamaysky: Columbia Business School, hm2646@columbia.edu. Shen: Columbia Business School, yshen21@gsb.columbia.edu. We thank Larry Glosten and Laura Veldkamp, as well as seminar participants at Boston University, Cornell, Columbia, the Rutgers Equilbrium Theory Workshop and the SIAM Conference on Financial Mathematics for valuable comments.

and prominent investors are not simply passive streams of information: the level of information they provide depends on investor demand. We then find that asset prices can change dramatically in response to changes in the supply and demand for information.

To capture these ideas, we develop a dynamic model of information and asset prices in which investor information choices influence the level of available public and private information about fundamentals. We study two types of feedback. In the first, as more investors become informed, more information about fundamentals becomes available; this mechanism magnifies the asymmetry between informed and uninformed investors, it tends to increase price volatility, and it can amplify small shocks into large price drops. In the second type of feedback we study, an increase in the number of informed investors leads to greater public information through leakage or disclosures; this mechanism has a stabilizing effect. These two mechanisms reflect two responses to greater demand for information: an increase in potentially knowable information, and an increase in the fraction of information made public.

Before describing the model further, we provide some motivating examples. In 2009, the Greek government revised its estimated budget deficit. This revision triggered a large increase in investor demand for information about Greek debt, as reflected for example in media attention and internet searches. Figure 1 shows the large and persistent increase in the number of Bloomberg articles mentioning Greece starting in 2009, and it provides at least circumstantial evidence that greater demand for information was met with greater supply. More information followed in the form of further revisions to official statistics, revelations about falsified data, stories of investment banks complicit in masking true conditions, research reports by industry analysts and non-governmental organizations, and a downgrade to junk by Standard & Poor's. The price of Greece's debt dropped sharply as the volatility of its sovereign credit default swap spreads rose. Consistent with this scenario, our model will describe circumstances in which feedback between the demand and supply of information leads to large price drops and increased volatility.

In June 2007, Bear Stearns disclosed that two of its hedge funds were on the brink of failure, fueling investor demand for information about the type of subprime mortgages in which the funds had invested. Indeed, Gorton and Ordoñez (2014) and Dang, Gorton, and Holmström (2012, 2019) have argued that the demand for information about "safe" collateral triggered the ensuing crisis. In our narrative, as more investors chose to incur the cost of becoming informed, more information became available — through revised credit ratings, academic and industry research, and regulatory reports. Less informed investors, fearing an informational disadvantage, fled to safer assets. Abstracting from the specifics of this setting, we calibrate

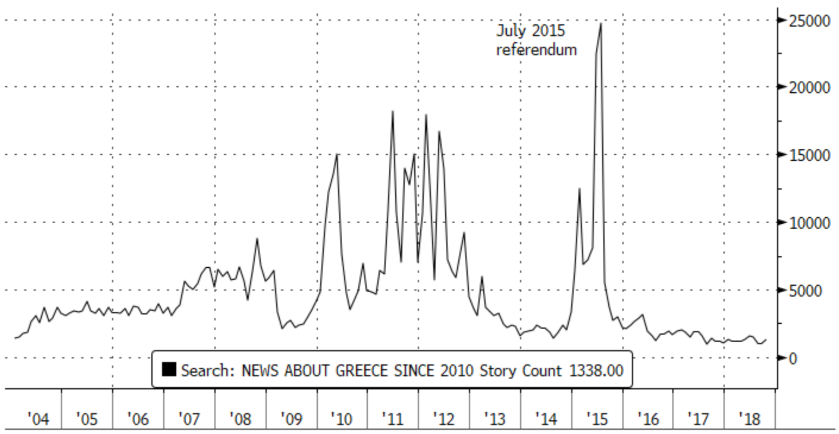

Figure 1: Number of Bloomberg articles per month that mention Greece.

our model to stock market data and find that the dynamics of information precision alone, even without negative news about fundamentals, can produce crisis-like effects.

The examples of Greek debt and subprime mortgages illustrate cases of destabilizing feedback in available information. We also study a stabilizing role for changes in the mix between public and private information. Here we turn to the literature on the information value of analysts' research. In particular, Kelly and Ljungqvist (2012) use brokerage closures as an exogenous shock to the level of coverage. Following this method, Balakrishnan et al. (2014) find that firms respond with greater information disclosure, and Chen et al. (2018) find that hedge funds respond with greater information acquisition. These are the types of feedback effects between shocks to information availability and investor information choices our model describes. We calibrate our model to this setting and find that it can reproduce the key empirical findings in Chen et al. (2018).

Motivated by these types of examples, our model combines exogenous shocks to the quality of available information, an endogenous response by investors who may choose to become informed at a cost, and feedback from investor information choices to the information environment. Information about fundamentals falls in three categories: publicly known, privately knowable at a cost, and completely unknowable. [1] In the interest of clarity, we treat separately the cases in which the share of public information and fraction of knowable information vary.[2] These two

[1]Publicly known information includes a product release that is covered in the New York Times. Privately knowable but costly information includes the performance of a firm's supplier network, which can be analyzed with painstaking analysis of public information. And information that is unknowable includes the outcome of a future medical trial relative to expectations.

[2]Treating both together requires working with a two-dimensional state space but does not present any conceptual or theoretical difficulties.

settings lead to two types of feedback, with the number of informed investors influencing the precision of either knowable information or specifically public information.

Examples of changes to the information environment include regulatory changes (e.g., the Sarbanes-Oxley Act or Regulation Fair Disclosure), changes in accounting standards, increased media attention to a topic, voluntary disclosure decisions by firms, and increased investment in research by financial intermediaries. In practice, these types of changes may influence both privately knowable and public information. Our model makes a sharp distinction between different types of information to clarify the distinct effects they can have.

In more detail, we develop an overlapping generations (OLG) model with a single risky asset, which pays a dividend each period, and a riskless asset. In each period, a new generation of investors observes the information environment, decides whether to become informed at a cost, sets optimal demands, and trades to clear an exogenous net supply of shares. Market clearing determines the price. At the end of the period, these investors receive their dividend and sell their shares at the new price. The notion of "generations" should not be taken literally in our setting; the OLG framework simply provides a tractable dynamic setting to model changes in information, and it ensures that investors care about future prices as well as the next dividend.

Crucially, in making their information choices at the start of the period, investors take into account the distribution of shocks to information precision and the feedback from information choices in the current period to future precision. The future precision will affect the end-of-period asset price and thus investors' capital gains. Incorporating such time variation in information precision into a rational expectations setting is technically challenging, and this is one of the key contributions of the paper. Using this framework, we show that information shocks can lead to large drops in prices and increases in volatility, and these effects are particularly noteworthy given that our investors have rational expectations.

The interplay between information and asset prices is often studied through single-period models of the type in Grossman and Stiglitz (1980), Hellwig (1980), Admati (1985), and a large subsequent literature. But there are several important features available in a dynamic model that are inaccessible in a single-period model, and these merit discussion. The first two important features we have already highlighted: feedback from investor information choices to information precision, and exogenous shocks to precision. Our emphasis is on feedback effects, and these cannot be captured in a single-period model. Persistent exogenous shocks are also important — they include the loss of analyst coverage discussed above, and also changes in accounting standards, like the introduction of mark-to-market accounting, and changes in regulatory policies on disclosures. In a single-period model, exogenous shocks are often approximated by changes in model parameters, but such changes are necessarily outside the model and, in particular, not contemplated by the agents in the model. In contrast, our agents' beliefs take into account that the economy can transition between different information regimes; such transitions are therefore a feature of the model itself.

A third important feature of a dynamic model is that it can capture two distinct aspects of an increase in available information: greater information reduces uncertainty about the next dividend but can *increase* volatility in future prices and thus in capital gains. The first of these effects is clear — the information we model is information about dividends. To appreciate the second effect, note that in the absence of dividend information, price volatility is driven entirely by supply volatility; but when some investors have dividend information, this information is partly reflected in the price, so volatility in the signal adds to volatility in the price. In a single-period model, the price merely determines the cost of a claim to an end-of-period dividend. With overlapping generations, investors earn the change in price over the period as well as a dividend, so the variance in this return affects their investment decisions at the beginning of the period. The two information effects, on dividends and an end-of-period prices, are potentially offsetting and lead to more complex tradeoffs than can be captured in a single-period setting.[3] We will see that this dual role of information in dynamic models can lead to starkly different conclusions than those of static models.

To the best of our knowledge, our model is the first to capture a stochastic information environment, endogenous investor information choices, and feedback from these choices to available information. Spiegel (1998) and Dutta and Nezlobin (2017) develop overlapping generations models in which all investors have the same information. Watanabe (2008) extends Spiegel's (1998) model by introducing asymmetric information. Biais, Bossaerts, and Spatt (2010) also model asymmetric information in an OLG setting. In their model, as in Watanabe's (2008), the fraction of informed investors and the precision of their signals are fixed and exogenous. Wang (1993) develops a continuous-time model of trading among differentially informed investors with a fixed fraction of informed investors and a fixed information environment; Wang (1994) is a discrete-time version of the model that investigates trading volume. In Avdis (2016), the fraction informed is endogenous but does not affect the information environment. The OLG model of Farboodi and Veldkamp (2017) incorporates a changing information environment, but the change is limited to a deterministic increase in investor information processing capacity over time. Signal precision also changes deterministically over time in Brennan and Cao (1997).

[3] Similar tradeoffs arise in the multiperiod models of Avdis (2016) and Dutta and Nezlobin (2017), but those models do not include feedback effects.

Through the feedback from information demand to information precision in our model, information shocks are amplified and can produce crisis-like dynamics or, less dramatically, business cycle fluctuations. Amplification in our setting arises purely through an information channel. This adds to other amplification mechanisms, such as financial frictions (as in Bernanke and Gertler 1989, Kiyotaki and Moore 1997, and Adrian and Shin 2010) or leverage (as in Lorenzoni 2008 and Bianchi 2011). We are not suggesting that other mechanisms are less important, but rather highlighting the role that information feedback alone can play.

Information revelation is at the center of the crisis explanation of Gorton and Ordoñez (2014). In their account, a crisis results when lenders choose to acquire information about borrowers' collateral; with less information available, borrowers with poor collateral have access to credit, and the increased supply of credit sustains higher growth. We work in an entirely different framework, but one contrast is particularly noteworthy. In Gorton and Ordoñez (2014), the information revealed is bad news; following an aggregate shock, some unobservable amount of collateral becomes bad, thus inducing more information acquisition. In our setting, only the precision of information changes — a shock may bring more news or less news, but not specifically good or bad news. An increase in precision leads to a price drop when it magnifies the information asymmetry between informed and uninformed investors, leading the uninformed to reduce their demand for the risky asset. Of course, a crisis is more likely to be precipitated by bad news, and adding a directional shock would likely further amplify the effects we observe; but our model isolates the role that the dynamics of information precision alone can play, without a negative shock to fundamentals.

To illustrate these effects, we calibrate our model to stock market data. The equilibrium dynamics of the calibrated model fluctuate between two regimes, one with low volatility and high prices, and one with high volatility and low prices. The model can spend long intervals in each regime. A transition from one to the other can be sudden and result in a price drop of 10%, with no change in fundamentals. The two regimes emerge from investor information choices; we do not impose them in setting up the model. Furthermore our investors are fully rational: they understand that the economy can transition from one regime to the other.

This pattern has important implications: in times of market stress, a policy change that makes more information available is potentially destabilizing. We emphasize "policy change" because the information environment in our model is persistent, so the effects we study go beyond a single announcement. Releasing positive information may help calm markets, but our model indicates that this effect must be weighed against the increased volatility that can accompany increased information. There is an extensive literature studying the downsides of increased disclosures, and these include reducing risk-sharing opportunities, distorting incentives for managers, inducing agents to underweight private information, and crowding out incentives for the production of additional information; see Goldstein and Yang (2017) for a survey. But the impact on prices and volatility we identify in our dynamic model is new in this context. Indeed, working in a single-period setting, Goldstein and Yang (2017) show that greater disclosure decreases return volatility, again highlighting the difference in perspective in a multiperiod model.

In the second version of our model, information dynamics can have a stabilizing rather than amplifying effect. In this formulation, as the number of informed investors increases, the amount of *public* information also increases, either because some of the informed leak their information or because a firm responds with increased disclosure. A large literature finds an association between reduced asymmetric information, higher liquidity, and a lower cost of capital (see, for example, Balakrishnan et al. 2014, Botosan 1997, and Ng 2011), which supports a motive for firms to disclose more public information in response to an increase in the number of privately informed investors. Chen, Matsumoto, and Rajgopal (2011) find that firms with lower institutional ownership are more likely to suspend guidance on earnings; put differently, this finding is qualitatively consistent with an increase in informed (institutional) investors driving an increase in public information.

Chen et al. (2018) study the effect of an exogenous loss of analyst coverage resulting from the closures or mergers of brokerage research departments. They document the following effects: price information efficiency falls; hedge funds trade more aggressively in the stock around earnings; hedge funds experience better investment performance in the stock; sophisticated investors increase their information acquisition; conditional on a large hedge fund presence in the stock, the loss of coverage has a smaller effect on price efficiency. We calibrate our model to this setting and show (in an appendix) that it reproduces these findings. We then use the calibrated model to examine the empirical results in Kelly and Ljungqvist (2012) about the stock price impact of brokerage closures. First, we show that our model generates a brokerage closure price impact whose sign and magnitude are consistent with the data. Second, the model-implied price impact from a brokerage closure is considerably smaller in a dynamic setting than it is in a static setting. Finally, introducing positive feedback from the informed to the fraction of publicly knowable information further diminishes the price effect.

We present our model in Section 2. Section 3 solves the model and states our main theoretical results. Section 4 studies changes in the level of knowable information and shows that feedback can lead to two information regimes, using parameters calibrated to market data. Section 5 contrasts our work with models of strategic complementarity in information acquisition. In an appendix, we study changes in the public-private mix of information, focusing on the application to the loss of analyst coverage. The appendices also provide proofs of our theoretical results. A Supplementary Appendix provides details of our calibration and numerical calculations.[4]

2 Model

2.1 Dividends and Information Dynamics

A single infinitely-lived security pays a dividend in each period. The dividend paid at the end of period t is given by

$$D_{t+1} = \bar{D} + \rho(D_t - \bar{D}) + M_{t+1} = \underbrace{(1-\rho)\bar{D}}_{\mu_D} + \rho D_t + M_{t+1}. \tag{1}$$

The innovation M_{t+1} decomposes as

$$M_{t+1} = \underbrace{m_t + \theta_t}_{\tilde{m}_t} + \epsilon_{t+1},$$

with the following interpretation: m_t is known to informed investors; θ_t is public information; $\tilde{m}_t$ is the knowable portion of the innovation; and ϵ_{t+1} is unknowable at the beginning of period t. These are mean zero, normally distributed random variables, independent across time,[5] with variances given by

$$\mathsf{var}(\tilde{m}_t) = f_t \mathsf{var}(M) \qquad \text{and} \qquad \mathsf{var}(\epsilon_{t+1}) = (1-f_t)\mathsf{var}(M). \tag{2}$$

and

$$\mathsf{var}(m_t) = \phi_t \mathsf{var}(\tilde{m}_t) \qquad \text{and} \qquad \mathsf{var}(\theta_t) = (1-\phi_t)\mathsf{var}(\tilde{m}_t). \tag{3}$$

Thus,[6]

$$\begin{aligned} f_t &= \text{fraction of dividend innovation that is knowable;} \\ 1-\phi_t &= \text{fraction of knowable part of dividend innovation that is public.} \end{aligned}$$

[4] Available at `https://sites.google.com/view/hmamaysky`.

[5] More precisely, they are conditionally independent given all (f_t, ϕ_t).

[6] Boot and Thakor (2001) distinguish three types of information disclosure. In our setting, these correspond roughly to an increase in $f_t\phi_t$, an increase in f_t, and a decrease in ϕ_t, respectively.

Before making investment decisions in period t, all agents observe θ_t and D_t, and the time-t informed agents observe m_t. A fraction $\lambda_t \in [0,1]$ of agents are informed at time t. The time-t uninformed agents, representing $1-\lambda_t$ of the population, in addition to observing θ_t and D_t, also observe the market clearing price P_t. Since the market-clearing price contains information about m_t through the demands of the informed traders, the uninformed also make rational inferences from the price about the innovation m_t. The price is not fully revealing about m_t because of the presence of unobservable supply shocks. In this respect, for a given f_t and ϕ_t, our information environment is the same as in Grossman and Stiglitz (1980).

The innovation of our paper is to allow the information environment, as represented by f_t and ϕ_t, to evolve over time in response to exogenous shocks and in response to information decisions made by past generations of investors.

Consider, first, the evolution of the precision f_t. We want f_t to be subject to i.i.d. shocks $\epsilon_{f,t+1}$, and we want the effects of these shocks to be persistent. The simplest model consistent with these objectives is an AR(1) process, $f_{t+1} = a_f + \kappa_f(f_t - a_f) + \epsilon_{f,t+1}$, where a_f and κ_f are constants. But we also want to allow the possibility that the precision f_t changes as the fraction informed λ_t changes. The simplest modification that captures this type of feedback is a model of the form

$$f_{t+1} = a_f + b_f\lambda_t + \kappa_f(f_t - a_f) + \epsilon_{f,t+1},$$

for some constant b_f. To be consistent with the interpretation of f_t as a measure of signal precision in (2), we need to restrict f_t to values between 0 and 1. We therefore apply a mapping $\Pi_{\mathcal{D}}$ to the right side of this equation, where $\Pi_{\mathcal{D}}$ maps the real line to a set $\mathcal{D} \subseteq [0,1]$.[7] We thus arrive at our first model of the information environment:

(**Level of knowable information.**) We set $\phi_t \equiv \phi_0$ and let f_t evolve according to

$$f_{t+1} = \Pi_{\mathcal{D}}\left(a_f + b_f\lambda_t + \kappa_f(f_t - a_f) + \epsilon_{f,t+1}\right). \tag{4}$$

Similar considerations lead us to our second model, in which the parameter ϕ_t varies:

(**Public-private mix.**) We fix $f_t \equiv f_0$ and let ϕ_t evolve according to[8]

$$\phi_{t+1} = \Pi_{\mathcal{D}}\left(\phi_t + b_\phi\lambda_t + \epsilon_{\phi,t+1}\right). \tag{5}$$

[7] In the simplest case, $\Pi_{\mathcal{D}}(x) = \min(1, \max(0,x))$ projects x to $[0,1]$. For some of our theoretical results in Section 3 and for our numerical results, we will discretize f_t and ϕ_t to finite subsets of the unit interval, but for now we keep the discussion general.

[8] Our results go through if ϕ_t is mean-reverting, but in our main application we will want to consider a lasting loss of information, so we use (5) from the outset.

Time t

First, agents observe f_t, ϕ_t and decide whether to become informed at cost c_I

Then, simultaneously,

- Agents observe public information P_t, D_t, θ_t
- Informed observe private information m_t
- Agents submit their demands q_t^I, q_t^U
- Price P_t clears the market,

$$\lambda_t q_t^I + (1 - \lambda_t) q_t^U = \bar{X} + X_t$$

Time $t+1$

Agents sell back their shares at price P_{t+1}

Agents earn excess return

$$P_{t+1} + D_{t+1} - RP_t$$

Figure 2: Sequence of events in each period.

The specifications in (4) and (5) provide the simplest models that capture persistent, stochastic time variation in the information environment and, most importantly, feedback from the fraction informed λ_t to the available information. We emphasize these qualitative features over the specific functional forms in (4) and (5).[9] The signs of b_f and b_ϕ, which we discuss momentarily, determine the directionality of the feedback effect.

In equilibrium, the state of the economy will be either f_t or ϕ_t, depending on which version of the model we consider. We could model the joint evolution of f_t and ϕ_t, but we keep them separate to emphasize different consequences of the two settings. (Notice that (4) and (5) are not mutually exclusive, as the precision of private information and the mix of public and private information could change simultaneously.) To cover both versions of the model, we will often refer to (f_t, ϕ_t) as the information state, though one of the two variables is always fixed.

Similarly, we will sometimes write model quantities (such as price coefficients) as functions of (f_t, ϕ_t).[10] In particular, we restrict the fraction informed λ_t to be a function $\lambda(f_t, \phi_t)$ of the information state. This means that $\lambda_t = \lambda(f_t)$ in (4) and $\lambda_t = \lambda(\phi_t)$ in (5). The specific form of $\lambda(\cdot)$ will ultimately be determined endogenously as investors make their information choices.

Figure 2 summarizes the timing of the model. In each period, investors first observe the

[9] An interesting topic for future research is to derive the relationships in (4) and (5) from first principles. However, the main modeling challenge in our paper is to embed the information dynamics of (4) and (5) in a rational expectations framework (see Section 3). Therefore such a micro-foundation is beyond the scope of the present paper though we comment on some potential mechanisms in Section 2.1.1.

[10] When f is restricted to a finite set, a function of f can be represented as a vector indexed by the values of f, and similarly for ϕ. We use this representation in our numerical calculations.

information state and choose whether to become informed at a cost c_I. Investors then observe public information and set their demands as functions of the price, which determines the price through market clearing. At the end of the period, investors sell their shares at the new price, and the process repeats.

2.1.1 Discussion of Information Dynamics

Taking $b_f > 0$ in (4) implies positive feedback from the fraction informed λ_t to the precision of knowable information. Taking $b_\phi < 0$ in (5) implies leakage of information (or increased disclosure) as the fraction informed increases. We should emphasize that with this sign convention the feedback from informed investors to the information environment in both (4) and (5) goes in the *same* direction: informed investors increase information availability in the next period. In (4), $b_f > 0$ implies that a larger number of informed investors makes more information knowable in the next period. And in (5), $b_\phi < 0$ implies that a larger number of informed investors makes more of the knowable information in the next period publicly available.[11] A scenario in which greater demand for information results in lower supply does not seem plausible, so we assume $b_f > 0$ and $b_\phi < 0$ throughout.[12]

We noted in the introduction that changes in technology and policies can contribute to changes in the information environment; as more specific examples, here we add that Rösch, Subrahmanyam, and van Dijk (2017) document time variation in measures of market efficiency, Brancati and Macchiavelli (2019) find time variation in information precision measured by dispersion in analyst forecasts, and Glasserman, Li, and Mamaysky (2019) find significant time variation in the informativeness of news articles from Thomson-Reuters. As additional considerations supporting feedback effects, we note that Gentzkow and Shapiro (2010) find that newspapers slant their political news to cater to subscribers; it is natural to expect that the resources devoted to business news would also respond to subscriber interest. Brancati and Macchiavelli (2019) report increased analyst coverage of banks at the onset of the financial crisis. As emphasized by Veldkamp (2006), the production of news is characterized by high fixed costs and low marginal costs; with this combination, increased demand results in lower prices. Equation (4) captures a similar response if we suppose that information suppliers respond to an increase in demand λ_t by improving precision f_t rather than lowering prices. Glasserman, Li, and Mamaysky (2019) find a positive correlation between news informativeness and the level of

[11] The opposite signs of b_f and b_ϕ reflect that f_t measures the knowable portion of information and ϕ_t measures the portion of that information that is *private*.

[12] Our model can, of course, be solved under the alternative assumption that $b_f < 0$ and $b_\phi > 0$. However, We do not consider this case to be empirically relevant.

intermediary capital; this correlation is consistent with (4) in the sense that higher levels of intermediary capital should be associated with higher levels of λ_t because financial intermediaries act as informed investors and providers of investing information. Frankel, Kothari, and Weber (2006) find that analyst forecasts are more informative for more volatile stocks; this pattern provides some indirect support for (4) to the extent that a larger λ_t is associated with greater price volatility, a question we return to later.

The accounting literature documents many related associations which suggest feedback consistent with (5). Botosan (1997) finds that greater disclosure leads to a lower cost of capital when analyst coverage is low, suggesting that a loss of coverage should lead firms to provide greater disclosure. Healy, Hutton, and Palepu (1999) find that greater disclosure is associated with greater institutional ownership, which can be seen as a proxy for the number of informed investors in a stock. Chen, Matsumoto, and Rajgopal (2011) find that lower institutional ownership leads to lower public disclosures in the form of earnings guidance.[13]

As these examples make clear, the preponderance of the literature supports the dynamics in (4) and (5) along with the assumptions that $b_f > 0$ and $b_\phi < 0$.

2.2 Investor Optimization Problem

At the beginning of period t, a unit mass of new (young) investors enter the market, each endowed with wealth W_t, known at time t. For an investor who buys q shares of the risky asset at price P_t at the beginning of the period and sells the shares at the end of the period at price P_{t+1}, terminal wealth is given by

$$\begin{aligned} W_{t+1} &= R(W_t - qP_t) + q(D_{t+1} + P_{t+1}) \\ &= RW_t + q(D_{t+1} + P_{t+1} - RP_t), \end{aligned} \tag{6}$$

where $R > 1$ is the gross return on a riskless asset. It will be convenient to define the per period *net profit* from owning a single share of the stock as

$$\pi_{t+1} \equiv D_{t+1} + P_{t+1} - RP_t, \tag{7}$$

in which case the budget constraint becomes $W_{t+1} = RW_t + q\pi_{t+1}$. Agents who enter at time t consume their wealth at $t+1$ and leave the market. These agents set their demands for shares

[13] Anecdotally, the case of Herbalife is consistent with (5). Following claims by the prominent investor Bill Ackman that the company was a pyramid scheme, investor demand for information increased and the company responded by disclosing details of its operations.

of the risky asset at time t by solving

$$J_t^\iota \equiv \max_q \mathsf{E}\left[\mathsf{E}[W_{t+1}|\mathcal{I}_t^\iota, f_{t+1}, \phi_{t+1}] - \frac{\gamma}{2}\mathsf{var}(W_{t+1}|\mathcal{I}_t^\iota, f_{t+1}, \phi_{t+1})\Big|\mathcal{I}_t^\iota\right], \quad \iota \in \{I, U\}, \tag{8}$$

where $\mathcal{I}_t^U = \{f_t, \phi_t, \lambda_t, D_t, \theta_t, P_t, W_t\}$ is the uninformed agents' information set at time t, $\mathcal{I}_t^I = \mathcal{I}_t^U \cup \{m_t\}$ is the informed agents' information set, and $\gamma > 0$ is a risk aversion parameter. Similar objectives are used in Peress (2010), Van Nieuwerburgh and Veldkamp (2014), and Mondria (2010), and can be interpreted as expressing a preference for early resolution of uncertainty, in the sense of Kreps and Porteus (1978). Maximizing (8) is equivalent to maximizing

$$\mathsf{E}\left[v\left\{\mathsf{E}\left[-\exp(-\gamma W_{t+1})|\mathcal{I}_t^\iota, f_{t+1}, \phi_{t+1}\right]\right\}|\mathcal{I}_t^\iota\right],$$

with $v(u) = -\frac{1}{\gamma}\log(-u)$, if W_{t+1} is conditionally normal, as it will be in our equilibrium. We could allow investors to condition on past values of variables in their information sets in (8), but past information will be irrelevant, given our independence assumptions. We will use the notation $\mathsf{E}_t[\cdot] \equiv \mathsf{E}[\cdot|\mathcal{I}_t^U]$ to denote conditioning on the time t common information set.

In addition to investor demands for shares of the risky asset, we need to specify the supply. As in the OLG model of Allen, Morris, and Shin (2006), we assume that the supply X_t of the risky asset is independent and identically distributed from one period to the next. As explained in Allen et al. (2006), i.i.d. supply can be interpreted as the result of trading by price-insensitive noise traders who reverse their trades at the end of each period. New investors each period thus only clear a new exogenous supply shock.[14] We assume each X_t is normally distributed with mean zero and variance σ_X^2. Furthermore, we assume that there exists a positive net supply $\bar{X}$ of the risky asset, and that this fixed supply is constant over time.

2.3 Equilibrium

Given a function $\lambda : [0,1] \mapsto [0,1]$ (yielding the fraction informed $\lambda(f_t)$ or $\lambda(\phi_t)$), a *market equilibrium* is defined by a price process P_t and demands q_t^I and q_t^U, depending on the price and other time-t information $\mathcal{I}_t^I$ and $\mathcal{I}_t^U$, that clear the market,

$$\lambda_t q_t^I + (1 - \lambda_t) q_t^U = \bar{X} + X_t, \tag{9}$$

and for which q_t^ι solve (8), $\iota \in \{I, U\}$, for all t.

[14]Our model extends easily to allow persistent supply shocks, at the expense of adding an additional state variable, which complicates our numerical examples. See Avdis (2016) for a model in which supply persistence influences investors' decisions to become informed.

Market clearing and investor optimality define a market equilibrium, given a function λ that determines the fraction of investors who are informed. Next we define what it means for this fraction to be determined endogenously. As in our discussion of Figure 2, we suppose that investors at the beginning of the period can choose to become informed at a cost c_I, incurred at the beginning of the period but after observing the current information state (f_t, ϕ_t). Investors' decisions to become informed or remain uninformed thus define a mapping from the information state to the fraction informed, which is precisely λ. We will use the following:

Definition 2.1 (Endogenous fraction informed) *In the case of the f_t model (4), we call λ the endogenous fraction informed if it satisfies the following conditions for each $f \in [0,1]$:*

(i) $\lambda(f) = 0$ *and* $\mathsf{E}[J_t^I - Rc_I | f_t = f] < \mathsf{E}[J_t^U | f_t = f]$*; or*

(ii) $0 \leq \lambda(f) \leq 1$ *and* $\mathsf{E}[J_t^I - Rc_I | f_t = f] = \mathsf{E}[J_t^U | f_t = f]$*; or*

(iii) $\lambda(f) = 1$ *and* $\mathsf{E}[J_t^I - Rc_I | f_t = f] > \mathsf{E}[J_t^U | f_t = f]$*.*

In the case of model (5), we require the corresponding conditions with ϕ in place of f.

In case (ii), the fraction $\lambda(f)$ is the point at which the marginal investor is indifferent between becoming informed and remaining uninformed. Cases (i) and (iii) cover the possibility that one choice dominates the other and is therefore selected by all investors.

3 Model Solution and Variance Beliefs

The main challenge in our model is to combine a time varying information environment with agents' rational expectations. In this section, we state the main results on the solution of the model. We show that for arbitrary fixed $\lambda(\cdot)$ and what we call variance beliefs, the model admits a market equilibrium. Then, again for fixed $\lambda(\cdot)$, we give conditions for the existence of *self-consistent* variance beliefs and thus a rational expectations market equilibrium. Given a market equilibrium, we give conditions for an endogenous fraction informed $\lambda(\cdot)$. With this endogenous $\lambda(\cdot)$, it is possible that the initial variance beliefs are not self-consistent. Therefore, we combine results to give conditions for an *information equilibrium*, in which conditions for a rational expectations market equilibrium and an endogenous fraction informed are jointly satisfied.

For some of the results in this section (Propositions 3.2–3.4), and for numerical calculations, we discretize the state space by restricting $\mathcal{D}$ — the set of values that f_t and ϕ_t can take — to be a finite subset of $[0, 1]$. This discretization allows us to represent any function of f or ϕ as an n-dimensional vector. Our numerical procedure is discussed in the Supplementary Appendix.

3.1 Market Equilibrium

Proceeding with the first of these statements, we will show that, for any choice of λ, the model admits a market equilibrium in which the price process takes the form

$$P_t = a_t + b_t m_t + g\theta_t - c_t X_t + dD_t, \tag{10}$$

where g, and d are constants, and a_t, b_t, c_t are functions of the information state (f_t, ϕ_t) but do not otherwise depend on t.

To characterize investor demands, we will initially solve a more general version of their optimization problem (8), in which we do not assume that investors know the coefficients of the price process (the meaning of this will be clear momentarily). We then show how this leads to a market equilibrium.

If prices are given by (10), we can write terminal wealth W_{t+1} in (6) as

$$\begin{aligned} W_{t+1} &= RW_t + q(1+d)D_{t+1} + q(P_{t+1} - dD_{t+1} - RP_t) & (11)\\ &= RW_t + q\Big[(1+d)D_{t+1} + a_{t+1} + b_{t+1}m_{t+1} + g\theta_{t+1} - c_{t+1}X_{t+1} - RP_t\Big]. & (12)\end{aligned}$$

Note that m_{t+1}, θ_{t+1} and X_{t+1} are independent of D_{t+1}, and of any time t information. With a view to solving (8), we evaluate the conditional mean of terminal wealth as

$$\mathsf{E}[W_{t+1}|\mathcal{I}_t^\iota, f_{t+1}, \phi_{t+1}] = q\big[(1+d)(\mu_D + \rho D_t + \theta_t + \mathsf{E}[m_t|\mathcal{I}_t^\iota]) + a(f_{t+1}, \phi_{t+1}) - RP_t\big] + RW_t, \quad \iota \in \{I, U\}. \tag{13}$$

In the above, we write a_t from (10) as $a(f, \phi)$ to make explicit its dependence on the state variables, and discuss it further below. For the conditional variance, we use (11)–(12) to write

$$\begin{aligned} &\mathsf{var}(W_{t+1}|\mathcal{I}_t^\iota, f_{t+1}, \phi_{t+1}) \\ &= q^2(1+d)^2\mathsf{var}[D_{t+1}|\mathcal{I}_t^\iota, f_{t+1}, \phi_{t+1}] + q^2\mathsf{var}[P_{t+1} - dD_{t+1}|\mathcal{I}_t^\iota, f_{t+1}, \phi_{t+1}] & (14)\\ &= q^2(1+d)^2\left[\mathsf{var}(m_t|\mathcal{I}_t^\iota) + (1-f_t)\sigma_M^2\right] + q^2 V_B(f_{t+1}, \phi_{t+1}). & (15)\end{aligned}$$

In the second equality, we have used (1)–(2) and introduced the *variance belief* function V_B. If prices are indeed given by (10), then the "correct" (rational expectations) belief is given by

$$V_B(f,\phi) = b(f,\phi)^2 \phi f \sigma_M^2 + g^2(1-\phi) f \sigma_M^2 + c(f,\phi)^2 \sigma_X^2 \qquad \forall f, \phi, \tag{16}$$

as can be seen by comparing the last term in (14) and (15). However, we initially allow investors to have an arbitrary, strictly positive variance belief V_B, which is shared by all investors.

With arbitrary V_B, we do not have equality in (15); instead, we posit that investors solve their optimization problems (8) as though (15) held. In other words, investors solve (8) but with the conditional variance replaced by the right side of (15). Furthermore, we show in Appendix A.2 that given V_B, $a(f,\phi)$ in (13) is fully determined by the investor's optimization problem.

A market equilibrium with variance belief V_B is then a price process and investor demand functions that clear the market and solve (8) with this modification.

Proposition 3.1 *Under the f_t model (4) or the ϕ_t model (5), for any variance belief function $V_B(\cdot)$ bounded above and bounded away from zero, and any $\lambda(\cdot)$, there exists a market equilibrium with a price process of the form (10) in which a_t, b_t, and c_t are functions of the information state and do not otherwise depend on t.*

The proof of this proposition is in Appendix A.1. The market equilibrium of Proposition 3.1 is not in general a rational expectations equilibrium because the variance belief V_B may not coincide with the conditional variance $\mathsf{var}[P_{t+1} - dD_{t+1} | \mathcal{I}_t, f_{t+1}, \phi_{t+1}]$ in (14). But we can think of agents in the model as learning over time. Starting from an initial belief, investors set their demands and clear the market at a price of the form in (10). They (or the next generation) then observe the realized variance given by the right side of (16). They update their beliefs by setting V_B equal to this realized variance, and the process repeats. This in fact is how we solve our model numerically.

Given price coefficient functions b and c, the belief updating equation (16) defines a new V_B, and given V_B, Proposition 3.1 defines new coefficients b and c. (The coefficient a depends on V_B but does not enter in the update of V_B.) Combining the two steps yields a mapping from an initial pair of functions (b,c) to an updated pair (b,c). We have self-consistent beliefs, i.e., a *rational expectations market equilibrium*, at a fixed point of this mapping.

We prove the existence of a fixed point in the appendix under mild restrictions on model parameters. For technical reasons, in this analysis we limit the values of f_t and ϕ_t to a finite (but arbitrarily large) subset $\mathcal{D}$ of the unit interval. Appendix A.3 gives our most general parameter restrictions. For simplicity, here we state a special case that is easy to verify and holds in our numerical examples.

Proposition 3.2 *Suppose that $R \in [1, 1.2]$ and the model parameters satisfy*

$$(1 + d)\gamma\sigma_M\sigma_X \leq 0.28. \tag{17}$$

Then for any fixed $\lambda(\cdot)$, there exists a fixed point of the variance belief updating mapping. This fixed point defines a self-consistent variance belief and thus a rational expectations market equilibrium with prices of the form in (10).

The point of condition (17) is that we need $(1+d)\gamma\sigma_M\sigma_X$ to be small (see in particular footnote 15), and with the bounds on R we can show that 0.28 is small enough. See Appendix A.3 for more general conditions and the proof of the result.

Two special cases of Proposition 3.2 are worth mentioning. If we fix $f_t \equiv 0$ and $\lambda \equiv 0$ we have an OLG model without asymmetric information, similar to the one in Spiegel (1998). As in Spiegel's (1998) model, the coefficients in the price function can be expressed through solutions of quadratic equations.[15] If we fix f_t and λ at constant strictly positive values, we get a model similar to Watanabe's (2008), which has asymmetric information but a fixed information environment and no feedback.

3.2 Information Equilibrium

Propositions 3.1 and 3.2 take $\lambda(\cdot)$ as exogenous. We need to show that our notion of the endogenous fraction informed in Definition 2.1 is meaningful. Given a variance belief V_B, for every ϕ and f, we need to find a λ that makes investors exactly indifferent between paying the cost c_I of becoming informed or staying uninformed; if no such λ exists, we set λ equal to zero or one according to Definition 2.1. The details of this calculation are given in Appendix B.1.

The following proposition shows that this procedure does indeed generate an endogenous fraction informed. Because changing λ changes the evolution of f_t or ϕ_t, we need to be a bit more explicit about how we map these variables to the finite set $\mathcal{D}$ in (4) and (5); these details are discussed in the proof of the following proposition in Appendix B.2.

[15] In this case, the equation for a self-consistent variance belief reduces to solving a quadratic equation with two real roots, which is given by:

$$V_B^2 + \left[2(1+d)^2\sigma_M^2 - \frac{R^2}{\gamma^2\sigma_X^2}\right]V_B + (1+d)^4\sigma_M^4 = 0.$$

The two roots describe two market equilibria, one with high price variance and one with low price variance, and we need an upper bound on the left side of (17) to ensure that both roots are positive. However, the high variance equilibrium is unstable under arbitrarily small parameter perturbations; only the low variance equilibrium is robust to such changes. In our numerical experiments, we find that if we start from a low value of $V_B(\cdot)$ we converge to the low variance equilibrium.

Proposition 3.3 *Suppose the shocks $\epsilon_{f,t}$, $\epsilon_{\phi,t}$ have a density. Then for any strictly positive variance belief there exists an endogenous λ in the sense of Definition 2.1.*

This result holds, in particular, at a self-consistent variance belief. That is, given an exogenous $\lambda(\cdot)$ and the associated self-consistent variance belief V_B, we can solve for a new endogenous $\lambda(\cdot)$. However, once we change λ, the variance belief may no longer be self-consistent. When we solve the model with endogenous beliefs numerically, we start with an arbitrary[16] variance belief, we then solve for the endogenous fraction informed (as provided by Proposition 3.3), we then calculate the realized variance (16) using the endogenous λ, update the variance belief and repeat. We can formulate this process as starting with a pair of coefficient functions (b, c), from which we calculate λ and then a new (b, c). Combining the two steps yields a mapping from an initial (b, c, λ) to a new (b, c, λ). A fixed-point of this mapping defines an information equilibrium, in the sense that it yields a market equilibrium in which investors do not want to deviate from their information choices.

Existence of an Information Equilibrium

In this section, which is technical and can be skipped on a first reading, we establish that the procedure discussed in the prior paragraph has a fixed point. Proposition 3.3 ensures the existence of an endogenous λ, but it does not guarantee uniqueness, so we need a somewhat more general formulation to establish existence of an information equilibrium. For any coefficient functions (b, c), let $\Lambda_o(b, c)$ denote the set of λ satisfying Definition 2.1, which we know from Proposition 3.3 is nonempty. Let $\Lambda(b, c)$ denote the set of all convex combinations of elements of $\Lambda_o(b, c)$. If there is just one λ in $\Lambda_o(b, c)$, then $\Lambda(b, c) = \Lambda_o(b, c) = \{\lambda\}$. In our numerical experiments, instances of multiple λ satisfying Definition 2.1 occur rarely, and we have never encountered multiple solutions λ when using self-consistent variance beliefs. However, because we have not proved the uniqueness of λ, we need to work with the potentially larger set $\Lambda(b, c)$ in establishing the existence of an information equilibrium (b, c, λ).

A convex combination $\lambda \in \Lambda(b, c)$ represents a mixed equilibrium in the following heuristic sense. For each $f \in \mathcal{D}$, we can write

$$\lambda(f) = w_f \lambda_1(f) + (1 - w_f)\lambda_2(f), \tag{18}$$

with $w_f \in [0, 1]$ and $\lambda_1, \lambda_2 \in \Lambda_o(b, c)$, $\lambda_2(f) > \lambda_1(f)$. Interpret this to mean that a fraction

[16]More precisely, we start with a variance belief within the region where Proposition 3.2 ensures the existence of a fixed point.

w_f of investors thought equilibrium $\lambda_1(f)$ would be selected, and a fraction $1 - w_f$ thought $\lambda_2(f)$ would be selected. At the outcome $\lambda(f)$, the marginal investor is not indifferent between becoming informed or not. A fraction w_f of investors, expecting an outcome of $\lambda_1(f)$, will see $\lambda(f)$ as too high, and in response a fraction $w_f(\lambda(f) - \lambda_1(f))$ of investors will switch from informed to uninformed. Similarly, a fraction $1-w_f$, expecting $\lambda_2(f)$, will see $\lambda(f)$ as too low, resulting in a fraction $(1-w_f)(\lambda_2(f)-\lambda(f))$ switching from uninformed to informed. But then (18) implies that these effects offset each other, leaving the fraction informed at $\lambda(f)$.

We establish existence of an information equilibrium — a joint equilibrium in (b, c, λ) — within the broader class of information choices in $\Lambda(b, c)$. For the following, let $M(\lambda)$ be the set of market equilibrium parameters (b, c) consistent with the fraction informed function $\lambda = \{\lambda(f), f \in \mathcal{D}\}$; these are the fixed points in Proposition 3.2.

Proposition 3.4 *Suppose the conditions of Propositions 3.2 and 3.3 hold. Then there exists an information equilibrium* (b, c, λ)*, meaning that* $(b, c) \in M(\lambda)$ *and* $\lambda \in \Lambda(b, c)$*. In other words,* (b, c) *defines a market equilibrium given* λ*, and* λ *defines a (possibly mixed) information equilibrium given* (b, c)*.*

4 Price and Volatility Cycles

As discussed in the introduction, our model is motivated by the idea that as more investors become informed, more information may become available. This type of feedback can arise at the onset of market stress, when firms and governments are pressured to reveal more information in response to heightened investor attention. In this section, we will show that this dynamic can lead to periods of low and high volatility and high and low prices driven purely by changes in the information state, with no change in fundamentals. In other words, we can generate transitions similar to business cycles or even financial crises through changes in the level of information, without necessarily the release of negative information.

4.1 Dynamics of Information Precision

To provide insight into the model, we turn to a numerical example. We postpone details on parameter values to our discussion of model calibration in Section 4.2. The solid line in Figure 3 shows λ as a function of f in model (4). We calculate this curve by starting from a flat variance belief function and iteratively updating the variance belief and λ as discussed in Section 3. This iterative process converges very quickly in our numerical experiments.

At low levels of information precision f, the figure shows a flat section where $\lambda(f) = 0$; with little information available, no investor chooses to bear the cost of becoming informed. Once f increases to just above 0.4, we have a positive fraction of investors informed, and this fraction generally increases with the precision f.[17]

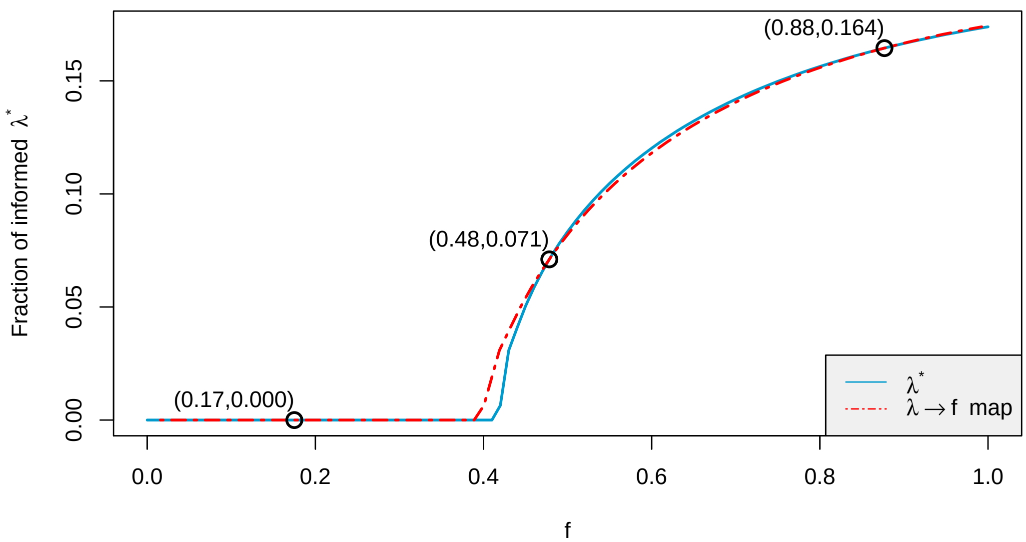

Figure 3: The solid line shows the fraction informed $\lambda(f_t)$ in information state f_t, and the dashed line shows the mapping from λ to f_{t+1} without exogenous shocks. Each circle shows a point where $f_t = f_{t+1}$ when the shocks in (4) are zero, labeled with its (f, λ) value.

To interpret the dashed line in Figure 3, we shut off the exogenous shocks in the evolution of f_t by setting $\epsilon_{f,t+1} \equiv 0$ in (4). The dashed line then shows the mapping from λ to the next value of f. That is, starting from any $f_t = f$ on the horizontal axis, reading up to the solid then across to the dashed line and back down to the horizontal axis yields f_{t+1}. Points where the two lines cross are equilibrium combinations of $(f, \lambda(f))$ in a model without exogenous shocks.[18]

Consider, for example, the circled point near $f = 0.48$, $\lambda(f) = 0.071$. Starting at that f, the endogenous fraction informed $\lambda(f)$ is precisely the value that keeps the information state at f under the evolution in (4) without endogenous shocks. The model still has feedback from λ to f (and f to λ), but f_t remains fixed. The same argument applies to the intersection near $f = 0.88$. In the lower left, the curves intersect throughout an interval where $\lambda(f) = 0$, and we have a fixed point at $(a_f, 0)$ because the dynamics in (4) drive f_t to a_f when $\lambda_t = \epsilon_{f,t+1} = 0$.

[17] For some parameter values, at f near 1 we have a small decline in $\lambda(f)$. The possibility of a decline in $\lambda(f)$ as f increases reflects the dual roles of information in a multiperiod model. Becoming informed benefits an investor by reducing uncertainty about the end-of-period dividend. However, as more investors become informed, the variance of the end-of-period asset price increases, so the net effect on the variance of an investor's end-of-period wealth is indeterminate.

[18] More precisely, the three circled points in the figure are cases where $f_t = f_{t+1}$ when $\epsilon_{f,t+1} = 0$.

But this perspective is somewhat misleading in a way that illustrates a difference between a genuinely dynamic model and a static one — a difference that will be important to the implications of the model. Without exogenous shocks, the three equilibria in the figure would seem to be equally valid solutions. But the middle equilibrium, near $f = 0.48$, is unstable: starting just to the right of the intersection will drive f_t to the equilibrium near $f = 0.88$, whereas starting just to the left will drive f_t to the interval where $\lambda(f) = 0$. The middle equilibrium is in some sense illusory, though it is a valid equilibrium without shocks.

If we reintroduce shocks in the evolution (4) and study the long-run distribution of f_t using the endogenous λ curve in the figure, we find that f_t spends significant time near $f = 0.175$, and it spends significant time near $f = 0.88$, but the region near $f = 0.48$ holds no particular attraction for the dynamic model with exogenous shocks. Figure 4 shows the steady-state f_t distribution (indicated by the blue circles in the left panel of the figure), calculated using a Markov chain representation.[19] The distribution is bimodal, showing that the economy spends the majority of its time in the vicinity of the two stable fixed points from Figure 3, but not near the middle, unstable fixed point in Figure 3.

The red triangles in the left panel of Figure 4 show the steady-state distribution of f_t when λ is held fixed at its equilibrium mean of 0.0731. The distribution is unimodal, indicating that f_t now spends most of its time near the middle of the interval (which happens to be in the vicinity of the unstable fixed point in Figure 3). This contrast points out the crucial role played by the endogeneity of λ and the feedback through $b_f \lambda$ (in equation 4) in generating two regimes. These regimes lead to price and volatility cycles in our model, as we discuss in the next section.

4.2 Model Calibration

In calibrating the model to the aggregate market, we take one period in the model to represent one month. We estimate a monthly dividend process of the form (1) using daily dividend data for the S&P 500 index from 1998–2018, then aggregating this up to the quarterly level (to mitigate seasonality effects), and estimating an ARMA(1,1) process for the quarterly dividend. From this we back out the monthly parameters $\rho = .967$ and $\sigma_M = 0.0471$. See the Supplementary Appendix for details.

We adopt the normalization $\bar{D} = 1$ and $\bar{X} = 1$, so dividends and share supplies are measured in units of their monthly averages. We calibrate σ_X^2 to match monthly turnover, meaning the number of shares traded per month divided by the shares outstanding. Recall from Section 2.2 that in each period t, investors buy the new supply X_t originating from liquidity demanders,

[19] See the Supplementary Appendix for details.

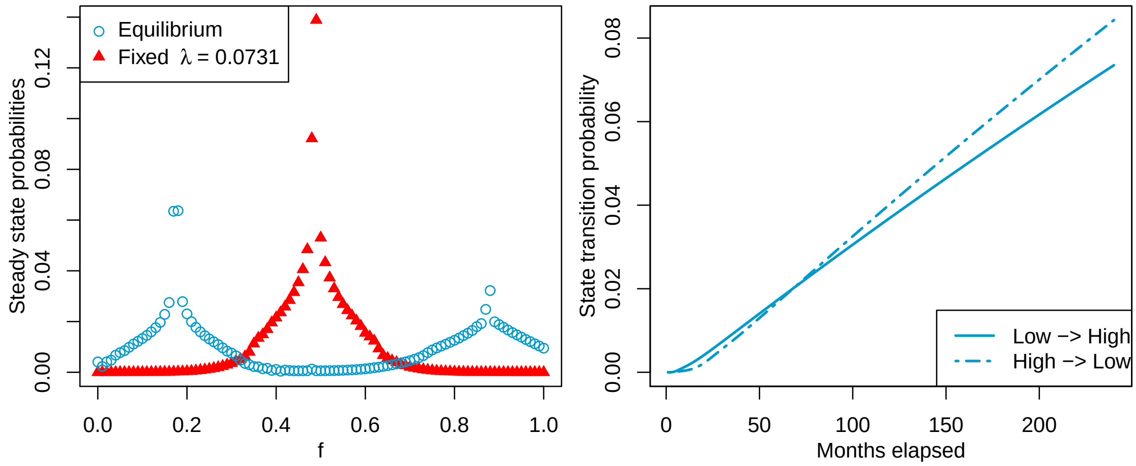

Figure 4: The left panel shows the equilibrium steady state distribution of f_t. The right panel shows the low-to-high regime transition probability $\mathsf{P}[f_{t+i} > 0.5|f_t = .17]$ (solid line) and the high-to-low regime transition probability $\mathsf{P}[f_{t+i} < 0.5|f_t = .88]$ (dashed line) as a function of i (measured in months).

and investors from the previous period sell back X_{t-1} shares to the previous period's liquidity demanders unwinding their trades. The total trading volume in period t is therefore $|X_t| + |X_{t-1}|$. Using the normality of the supply shocks, the expected volume per period becomes

$$\mathsf{E}[|X_t| + |X_{t-1}|] = 2\sigma_X \sqrt{\frac{2}{\pi}} \approx 1.596\sigma_X. \tag{19}$$

In the Supplementary Appendix, we find that the average weekly turnover of the Dow Industrials index is 0.065.[20] To model a period of stress, we assume that turnover, or the turnover expectation of market participants, is four times higher than normal, so $\sigma_X = 4 \times 0.065 \times 1/2 \times \sqrt{\pi/2} = 0.1629$.[21]

We use a monthly gross risk-free return of $R = 1.0015$ and set the risk-aversion parameter at $\gamma = 0.46$. This yields an annualized excess return of roughly 15%, which is not unreasonable for periods of stress. We choose a per month cost of being informed of $c_I = 0.2627$, which should be compared with a monthly aggregate average dividend of 1 (since $\bar{X} = \bar{D} = 1$). This high cost of information is comparable to the 2/20 fee structure of many hedge funds, and leads to an equilibrium number of informed of under 18% of the overall population.

[20] Lo and Wang (2000, Table 3) show that from 1987-1996, weekly turnover on a value-weighted index of NYSE and AMEX common shares was 1.25%. Therefore the monthly turnover on this index was $52/12 \times 1.25\% = 5.42\%$.

[21] The trading volume of SPY, which tracks the S&P 500 index and is one of the most liquid exchange-traded funds, has spiked by a factor of four during stress periods. For example, in 1/18/2011 the trailing month's average daily trading volume was 106.2 million and in 8/26/2011 the trailing month's daily trading volume was 407.0 million.

$\bar{D}$	$\bar{X}$	R	σ_X	ρ	σ_M	γ	c_I	ϕ_0	$\mathsf{P}[\epsilon_f \neq 0]$	ϵ_f	a_f	b_f	κ_f
1	1	1.0015	0.1629	0.967	0.0471	0.46	0.2627	0.35	0.03×2	0.135	.175	.384	.91

Table 1: Calibrated parameters for model (4)

For the dynamics of the f_t process in (4), we set $a_f = 0.175$, $\kappa_f = 0.91$ and $b_f = 0.384$. A value of κ_f close to 1 makes the information state persistent, and a positive b_f produces positive feedback from the fraction informed to the level of accessible information. When f_t is low, $\lambda(f_t) = 0$, and f_t is pulled toward a steady state level of $a_f = 0.175$.[22] We fix $\phi = 0.35$, implying that 65% of the knowable information is publicly known. This introduces a high degree of information asymmetry between informed and uninformed investors, and leads to a large drop in price in the high-information regime; see the discussion in Section 4.6.

Beyond these qualitative considerations, these specific parameters were chosen to produce plausible model dynamics. Finally, for the shocks $\epsilon_{f,t+1}$, we use a three-point distribution taking values $\{-0.135, 0, 0.135\}$ with probabilities $\{0.03, 0.94, 0.03\}$, so shocks are rare.[23] Model parameters are summarized in Table 1.

We should note that our results are robust to changes in model parameters. For a wide range of values in our non-dividend parameters (since ρ and σ_M are estimated from actual data) in Table 1, there exists a ϕ close to its base value of 0.35 which generates the bimodal f_t distribution and the large price drops that we discuss below. In fact, in many cases the resultant price drops are larger than the one that occurs under our base parameterization.[24]

4.3 Price Drops and Volatility Spikes

Figure 5 shows model quantities calculated using the parameters in Table 1. The first three panels show the price coefficient functions a, b, and c from (10). The lower right-hand panel shows the expected net profit from owning one share of the stock. When no investor is informed, no dividend information is reflected in the price, and $b = 0$. As f increases to the point where some investors become informed, b and c both increase, which drive up the price variance.[25] The

[22] Fama and French (2000) show the R^2's of year-ahead firm-level earnings forecasts of 5–20%.

[23] Parts of our theoretical analysis assume the shocks have a density to ensure continuity of the price coefficients in λ. A discrete distribution is simpler to work with numerically, and we achieve continuity by interpolating when we take expectations over state transitions of f, as discussed in the Supplementary Appendix. The interpolation has a similar effect as smoothing the shock distribution; moreover, a discrete distribution can be approximated arbitrarily closely by a density.

[24] This analysis is available from the authors.

[25] As b measures the sensitivity of the price to dividend information, the monotonicity of b parallels an empirical finding in Brancati and Macchiavelli (2019) that prices become more information-sensitive when information precision increases.

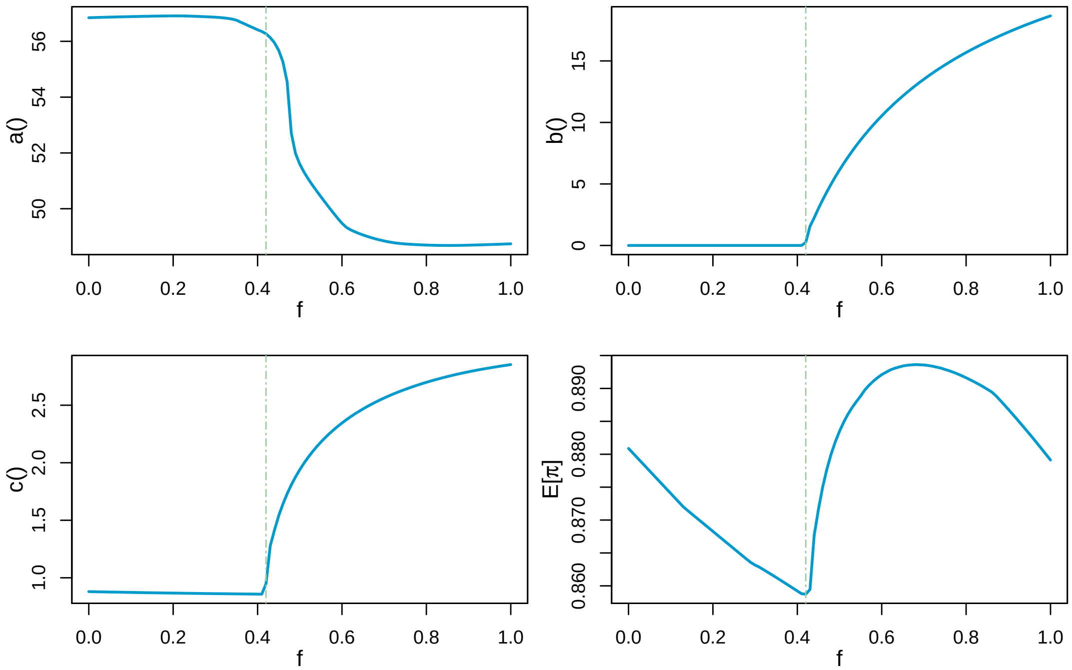

Figure 5: Equilibrium price coefficients $a(f)$, $b(f)$, and $c(f)$ and the expected investment net profit $\mathsf{E}[\pi]$ per period, using the parameters in Table 1. The vertical, dashed green line shows the point at which λ becomes positive.

increase in c reflects a higher compensation for accommodating supply shocks and is attributable to higher price variance and a growing informational disadvantage of the uninformed relative to the informed. Furthermore, as informed enter the market due to a higher f, a falls sharply, a phenomenon we will study in detail.

The left panel of Figure 6 shows the expected stock price $P_0 \equiv a(f)+d\bar{D}$. The price response is dramatic: a small increase in f leads to a price drop of 10%. This price drop results from an increase in information precision and an endogenous response of the number of informed investors. In stress times, more information revelation (and, as we will see in Section 4.6, greater information asymmetry) can destabilize the market. We will see in Section 4.4 that the conditional variances of net profit π in the right panel of the figure help explain the price drop.

Figure 7 shows an example of a simulated path of the model for 1,000 periods, or 83 years if each period is one month. The top two panels plot the evolution of the price P_t and variance[26] V_B, respectively. In this example, the market starts in the low volatility and high price regime, transitions to the high volatility and low price regime, and then transitions back. The figure

[26]The variance V_B in (16) increases monotonically in f, so the two variance regimes in Figure 7 correspond to two f regimes.

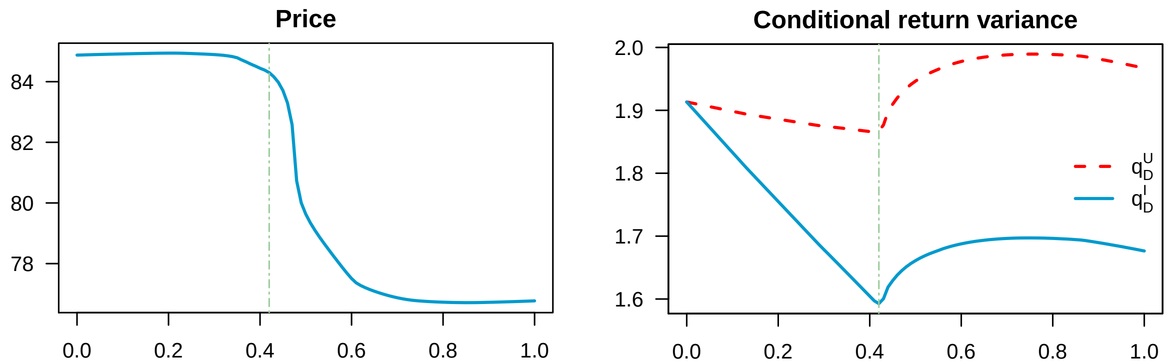

Figure 6: Equilibrium expected price $P_0 \equiv a(f) + d\bar{D}$ and conditional expectation of return variance for informed, q_D^I, and uninformed, q_D^U, investors, as functions of information state f.

shows a large drop in price associated with the spike in variance, resulting from an increase in f_t. This price drop is much larger than the within-regime price volatility, which is driven primarily by dividend fluctuations. The relationship between the levels of prices and variance across two regimes is further illustrated in the scatter plot at the bottom of Figure 7, which records the values from the top two panels.

It is customary to associate large declines in market values with the arrival of bad news. Following a 10% decline (the price drop in Figure 6) in an individual stock price or the overall market, one would expect media and expert accounts of what bit of bad news — a product failure, a CEO scandal, a change in government policy — triggered the fall. But in our setting it is simply more news — in the form of increased precision f_t — that drives investors, not necessarily good or bad news.

In the model of Gorton and Ordoñez (2014), the onset of a crisis is defined by the release of information about collateral quality.[27] But the additional information in Gorton and Ordoñez (2014) is negative information: collateral quality is inferior to what was previously believed. This revelation leads to either a collapse in lending or a diversion of productive resources to information acquisition, thus reducing growth in both cases. The mechanism in our model is entirely different and does not rely on adverse information about fundamentals.

In practice, an increase in the quantity of information without an accompanying positive or negative implication for fundamentals is rare, and this makes it difficult to disentangle a change in precision from a directional effect of news.[28] However, our model serves to isolate

[27]Their model, based on a very different framework, also exhibits an equilibrium with information cycles. In their setting, low information booms alternate with high information crashes.

[28]We discuss such an exogenous information shock in Section 5.2.

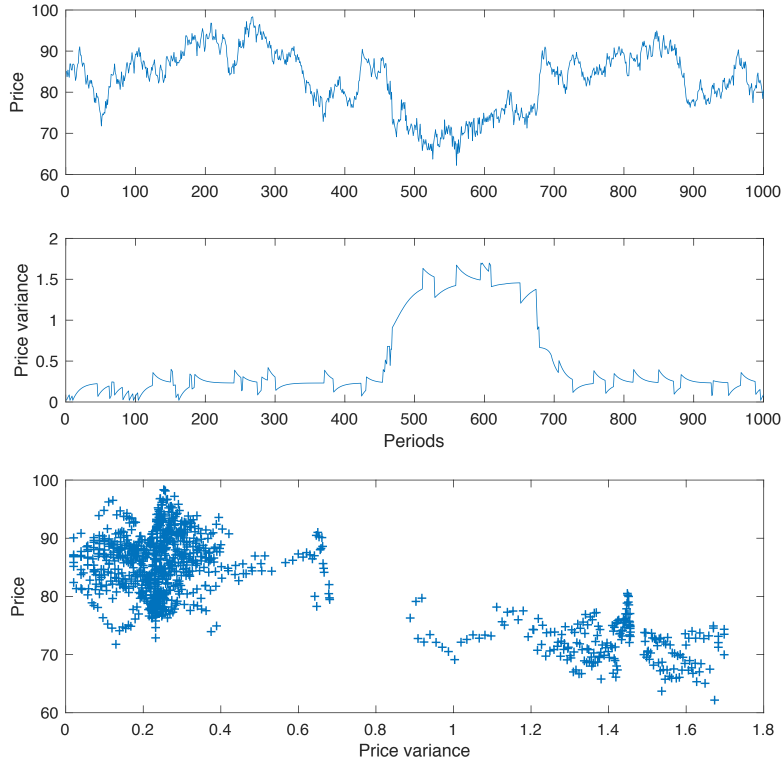

Figure 7: The top two charts show price P_t and price variance V_B in a simulated path of the model. The bottom chart reproduces the values in the top panels in a scatterplot.

the information component; our calibration indicates that this component alone can have a material effect. When the additional information is bad news as well as more precise news, we can expect the effects to be even greater.[29]

The potential for increased volatility from increased information has policy implications. A regulatory change that leads to persistently higher information precision for informed investors is potentially destabilizing in times of market stress.[30] Interestingly, in their analysis of disclosure of the results of regulatory stress tests for banks, Goldstein and Leitner (2018) conclude that disclosure is valuable only under adverse conditions. Our results do not conflict but rather reflect different considerations, as the objective in Goldstein and Leitner (2018) is optimal risk sharing among banks, and the information disclosed separates weak and strong banks.

The rest of this section explains the key features exhibited by our model, particularly the

[29]In a macro context, Cesa-Bianchi and Fernandez-Corugedo (2018) find that an increase in economic uncertainty results in a *decrease* in the risk premium, which is consistent with our results.

[30]The information disclosed about regulatory stress tests is disclosed publicly, but the design of scenarios and the interpretation of the results are technical matters that are arguably accessible only to informed investors who have acquired the necessary expertise.

large price drop illustrated in Figure 6 and the two regimes illustrated in Figure 7.

4.4 Decomposing Price and Volatility

The price drop in Figure 6 is driven by the drop in the $a()$ curve in Figure 5. In Appendix A.2 we show that a_t can be decomposed into two components,

$$a_t = \frac{(1+d)\mu_D}{R-1} - \frac{1}{R}\sum_{i=0}^{\infty}\frac{1}{R^i}\mathsf{E}_t[\pi_{t+1+i}], \tag{20}$$

where, π_{t+1+i} is the net profit from holding one share of the stock from $t+i$ to $t+1+i$, as in (7), and its conditional expectation is taken given f_t. From this expression, we see that the $a_t + dD_t$ component of P_t is the present value of all future expected dividend payments minus a discount reflecting the expected present value of all future net profits.[31] In the context of the Campbell (1991) and Vuolteenaho (2002) return variance decomposition, the second term in a_t represents the effect of a time-varying discount rate on the stock price. To understand how a change in information precision f_t creates a price drop, we need to understand the effect of f_t on the second term in a_t.

The stock's expected net profit over a single period is given by[32]

$$\mathsf{E}_t[\pi_{t+1}] = \underbrace{\gamma}_{\text{risk aversion}} \times \underbrace{\bar{X}}_{\text{asset supply}} \times \underbrace{\left(\lambda\frac{1}{q_D^I} + (1-\lambda)\frac{1}{q_D^U}\right)^{-1}}_{\text{average uncertainty}}. \tag{21}$$

The quantities q_D^I and q_D^U, defined in equations (25) and (26) of the appendix, represent the expected conditional variance of the net profit π, conditional on the information sets of the informed and uninformed investors, respectively. Equation (21) thus reflects the average return uncertainty faced by investors, weighted by the fractions of informed and uninformed in the economy, and scaled by $\gamma\bar{X}$.

The right panel of Figure 6 shows the expected conditional variances q_D^I (solid line) and q_D^U (dashed) of the net profit π for informed and uninformed investors. The shape of these curves reflects the tradeoff engendered by increased information precision. When f is low, an increase in f decreases the expected variance of net profits for informed and uninformed investors because more is known about next period's dividend, the D_{t+1} term in (7). As long

[31] If the second term in (20) is zero and the dividend equals its long-run mean, then $a_t + dD_t = \frac{(1+d)\mu_D}{R-1} + d\bar{D} = \bar{D}/(R-1)$.

[32] This is shown in equation (41) of the appendix. This expression generalizes the corresponding quantity derived from equation (A10) in Grossman and Stiglitz (1980).

as f is low enough so that $\lambda(f) = 0$ (to the left of the vertical, dashed line in the graphs) this is the only effect, and higher information precision lowers uncertainty. However, past the no-informed point, with f large enough that $\lambda(f) > 0$, the uncertainty of next period's net profit starts to increase, due to the increasing variance of P_{t+1} in the expression for π_{t+1} in (7). This effect outweighs the decrease in the variance of next period's dividend, and thus increases the conditional variances of the net profit. For high enough f the increased information about next period's dividend begins to dominate, and the expected conditional variance begins to fall again. This pattern depends crucially on the dynamic structure of our model: in a single-period setting, where investors care about the next dividend but not future prices, more precise information always reduces investment uncertainty.

Through (21), the common shape of q_D^I and q_D^U is inherited by $\mathsf{E}_t[\pi_{t+1}]$, as illustrated in the bottom-right panel of Figure 5. Recall from our discussion of Figures 3 and 4 that f_t spends most of its time near $f = 0.175$ or near $f = 0.88$. From the bottom-right panel of Figure 5, we see that $\mathsf{E}_t[\pi_{t+1}]$ is greater near $f = 0.88$ than it is near $f = 0.175$, indicating an increase in the expected profit from holding the stock as we move from the low-information regime to the high-information regime. This increase in expected profit is associated with a decrease in the current price of the stock, and it contributes to the price drop we see in Figure 6.

Notice, however, that the change in expected net profit across regimes is quite small, as indicated by the vertical scale in the lower-right panel of Figure 5. How does a small change in expected profit get amplified into a 10% price drop? The answer lies in the combination of the price discount reflected in (20) and the persistence of f_t.

The right panel of Figure 4 shows that transitions between $f_t \approx 0.175$ and $f_t \approx 0.88$ occur rarely. The chart shows the low-to-high transition probability $\mathsf{P}[f_{t+i} > 0.5|f_t = 0.175]$ (solid line) and the high-to-low transition probability $\mathsf{P}[f_{t+i} < 0.5|f_t = 0.88]$ (dashed line) as functions of the number of months i. Both probabilities grow very slowly, reaching only 6-8% even after 240 months, confirming that transitions between regimes are infrequent, and generating the bimodal stationary distribution in the left panel. As a consequence of this persistence, we expect the inequality $\mathsf{E}_t[\pi_{t+1}|f_t = 0.88] > \mathsf{E}_t[\pi_{t+1}|f_t = 0.175]$ (which we observed in Figure 5) to extend to $\mathsf{E}_t[\pi_{t+1+i}|f_t = 0.88] > \mathsf{E}_t[\pi_{t+1+i}|f_t = 0.175]$, for large i. Recall the present value of such terms is subtracted from the price P_t through the a_t coefficient in (20). Thus, even a relatively small single-period difference in expected profits around $f = 0.175$ and $f = 0.88$ is amplified to a large change in the price because the f_t process spends long periods in each of the two regimes before moving towards the other.[33]

[33]The same argument predicts a sharp decline in the $a()$ curve around the unstable fixed point near $f = 0.48$

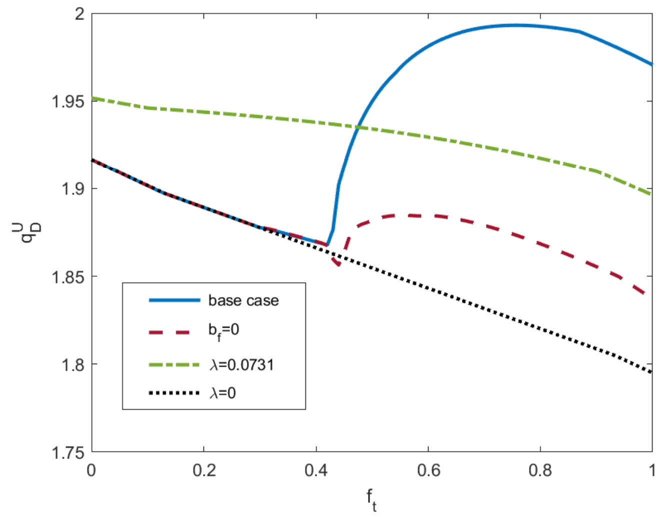

Figure 8: Conditional return variance q_D^U faced by uninformed investors in alternative models.

Are such infrequent regime transitions plausible? Barro (2009) estimates country level crises occur with a 1.7% per year probability. Assuming independence across time, a given country has a 29% (i.e., $1 - (1 - 0.017)^{20}$) probability of experiencing at least one crisis over a 20-year period. As we saw in Figure 4, the probability of a low to high state transition in our model is approximately 7% over a 20-year period. Our calibration therefore suggests that one out of four country-level crises may be accompanied by the information-driven price drop of our model. If crises are typically associated with positive shocks to f_t, the actual ratio may be higher.

4.5 Determinants of the Return Variance of the Uninformed

We saw in the previous section that large price drops in our model result from the combination of two properties: differences in the conditional variances q_D^I and q_D^U across f regimes, and persistence of these f regimes. In this section and the next, we take a closer look at the model features that drive these two properties.

Recall that the shape of the q_D^I and q_D^U curves in Figure 6 was our starting point for understanding the drop in the price that results from an increase in f. In particular, uninformed investors face greater conditional return variance q_D^U in the high-information regime near $f = 0.88$ than in the low-information regime near $f = 0.17$. Focusing on this property allows us to unpack the contributions of three key features of our model: the dynamic structure, the feedback from the fraction informed λ to the precision f_t, and the endogeneity of λ.

in Figure 3. Starting to the right of 0.48, f_t will tend to move toward 0.88, whereas starting to the left of 0.48, f_t will tend to move toward 0.175.

Figure 8 compares $q_D^U(f)$ in variants of our model. The solid line corresponds to our base case model and reproduces the dashed line in the right panel of Figure 6. Because our model extends over multiple periods (and as discussed in Section 4.3), an increase in information precision can increase the conditional return variance faced by the uninformed because the increase in price variance can outweigh the decrease in dividend uncertainty. The fact that q_D^U is larger in the high information regime than in the low information regime is a crucial factor in the price drop in Figure 6. The other curves in Figure 8 show the effect of shutting off the feedback effect (by setting $b_f = 0$ in (4)) or removing the endogeneity of the fraction informed (by setting $\lambda = 0$ or setting λ at its mean value of 0.0731). Each of these changes lowers the return variance at higher values of f_t and decreases the price discount in the high f_t regime. In short, the large price drop we see in Figure 6 relies on the combination of features that distinguish our model — dynamics, feedback, and endogenous information choices — and would be absent without this combination.[34]

4.6 The Role of Time-Varying Information Asymmetry

In this section, we show that information asymmetry plays an important role in generating price and volatility cycles in our model. In particular, large price drops occur when the economy transitions from low- to high-information asymmetry states. Whether this can happen is dictated by ϕ, the fraction of knowable information that is private; ϕ therefore plays a crucial role in the dynamic f_t version of our model. Figure 9 compares equilibrium $a()$ curves for different values of ϕ; the case $\phi = 0.35$ is the one we have analyzed thus far.

When $\phi = 0$ and all knowable information is public, the economy is characterized by no information asymmetry — the knowable information is equally known to all agents. The $a()$ curve corresponding to this no-asymmetry case is the highest one (shown as a solid line), indicating the smallest price discount relative to the present value of future dividends. The $a()$ curve in this case is quite insensitive to f_t. The $\phi = 1$ case represents the highest informational asymmetry possible in the model, and corresponds to the lowest $a()$ curve, representing a large price discount needed to induce the informationally disadvantaged uninformed agents to participate in risk sharing, regardless of the information state f_t. Only for intermediate values of ϕ can the economy transition from low- to high-information asymmetry states. Such regime

[34]The specific pattern of base case model can be interpreted as follows: As the endogenous λ increases with f, the endogeneity of λ and positive feedback $b_f > 0$ in dynamics (4) make f_{t+1} increase more quickly in f_t compared with the alternative models. Consequently, as f_t moves from low to high, the increase in next period's price variance outweighs the decrease in next period's dividend uncertainty, leading to a higher return variance in the high f regime.

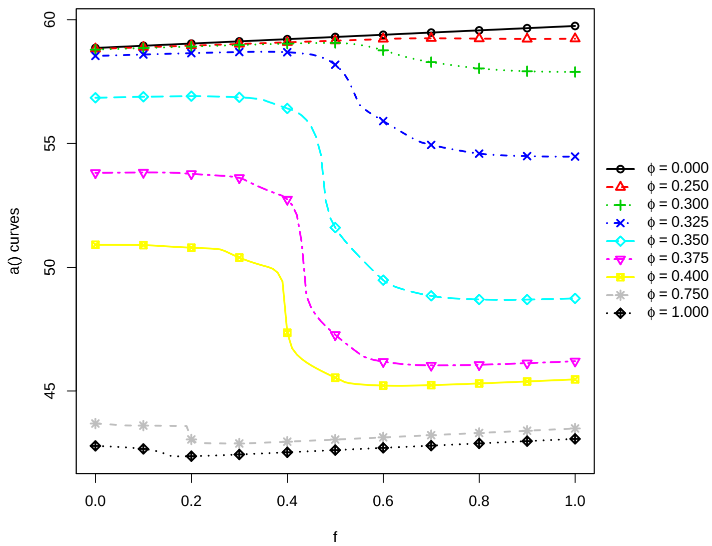

Figure 9: $a(\cdot)$ curves as functions of f for different values of ϕ

shifts are accompanied by large price changes.

The reason that prices in the case of $\phi \in \{0, 1\}$ do not change much across different values of f can be seen from Figure 10, which shows the steady-state distribution of f in the different ϕ models. When $\phi = 0$, there are no informed investors since all knowable information is public. With $\lambda = 0$ in (4), any positive $\epsilon_{f,t+1}$ shock quickly decays, pulling f_t back to its low-information fixed point. This dynamic is seen in the unimodal distribution, with the peak centered at a_f, when $\phi = 0$. Similarly, when $\phi = 1$, all knowable information is private, and λ is relatively large. Via the b_f term in the dynamics of f_{t+1} in (4), a relatively high λ produces a steady state distribution that is unimodal at $f = 1$. Any negative $\epsilon_{f,t+1}$ quickly dissipates as f is pulled back to one. In both cases, $\mathsf{E}_t[\pi_{t+1+i}|f_t = f]$ will be close to $\mathsf{E}_t[\pi_{t+1+i}|f_t =$ unimodal steady-state value$]$ for large i for any starting f, and a_t in (20) is consequently insensitive to f.

For intermediate values of ϕ (0.35 in our calibration), the equilibrium λ is in an intermediate range, and the tendencies of f_t towards a_f and towards the high-information fixed-point are balanced. The steady state f_t distribution becomes bimodal as can be seen in Figure 10. Therefore, a sequence of shocks can occasionally push the economy from one information regime to the other. And yet both regimes are very persistent. As in Section 4.4, this persistence amplifies differences in expected net profit $\mathsf{E}_t[\pi_{t+1}]$ at different f's to produce large price changes.

This effect results from an increase in information asymmetry, rather than from increase in information precision. When f_t is low, there is little information but also no information

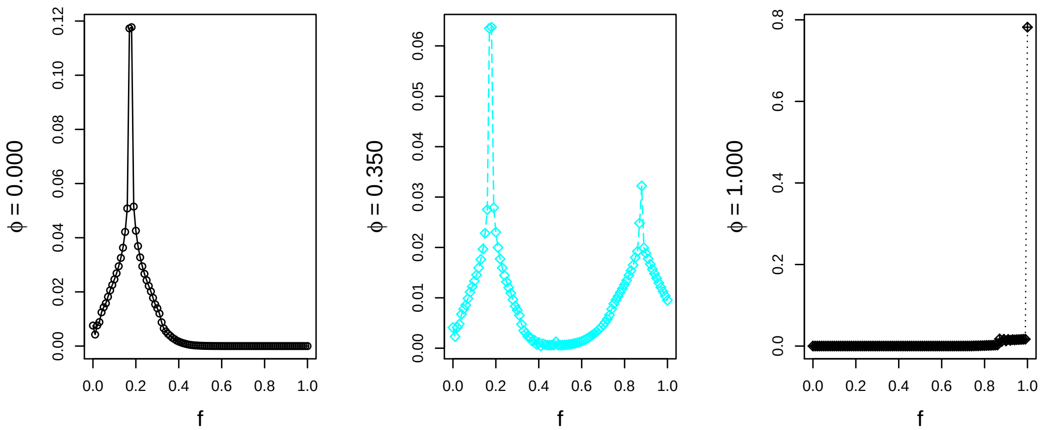

Figure 10: Steady-state distribution of f_t across different levels of ϕ.

asymmetry because all agents are uninformed. In this case, prices are high. But as f_t increases, private information becomes more revealing and some investors start to acquire it at a cost. The uninformed then find themselves at growing informational disadvantage and the price falls. As we argued in Section 4.5 this mechanism relies crucially on a dynamic information environment with feedback and endogenous information choices.

5 Strategic Complementarity and Substitutability in Information Choices

Models with multiple equilibria often exhibit some form of strategic complementarity. In models of information choice, the distinction between strategic complementarity and substitutability refers to whether the value of becoming informed increases or decreases as others become informed. In Section 5.1, we show that, perhaps surprisingly, the large price drop and the two regimes studied in Section 4 emerge without strategic complementarity. We then turn to model (5), in which ϕ evolves dynamically and f is fixed. This model, like the first, exhibits strategic substitutability in information choices. And yet, as we explain in Section 5.2, positive feedback in information production has a stabilizing effect in this setting. Taken together, these results show that the difference between the shock amplification we observe in the dynamic f model (4) and the stabilizing effect of greater disclosure in the dynamic ϕ model (5) is not driven by a contrast between complementarity and substitutability in information choices.

5.1 Dynamic Precision f_t

A key property of the Grossman and Stiglitz (1980) setting is that the value of becoming informed decreases as the number of informed investors increases. Subsequent work has investigated conditions in which the value of becoming informed increases as more investors become informed. Sources of this type of strategic complementarity identified in the literature include high fixed costs and low marginal costs in information production (Veldkamp 2006); certain deviations from normally distributed uncertainty (Chamley 2007); settings in which investors learn about supply as well as cash flows (Ganguli and Yang 2009 and Avdis 2016); other settings with multiple sources of information (Manzano and Vives 2011 and Goldstein and Yang 2015); and settings in which information acquisition affects cash flows (Dow et al. 2017). With few exceptions, these are static models, but they often result in multiple equilibria, with different asset prices in different equilibria.

In our dynamic setting, large price changes occur within the model, rather than through a change of equilibrium selection in a static model. But we will see that the contrast with earlier work goes beyond this feature. In order to explore strategic complementarity, we need to vary λ exogenously, shutting off its dependence on the information state f_t; at an (interior) endogenous λ, the marginal investor is, by definition, indifferent between becoming informed or remaining uninformed. With a fixed $\lambda > 0$ and $b_f > 0$, the dynamics in (4) tend to push the informativeness f_t to higher values at larger values of λ, which might suggest that the value of becoming informed increases with the fraction informed λ. But Figure 11 shows that this is not generally the case. Using parameter values from Table 1, the figure shows the expected utility gain $\mathsf{E}_t[J^I_{t+1} - J^U_{t+1}]$ from information acquisition as a function of f_t, at fixed levels of λ.[35] The difference $\mathsf{E}_t[J^I_{t+1} - J^U_{t+1}]$ measures the expected benefit of becoming informed in the next period, given the current information state.

Across a wide range of values of f_t, the curves are *decreasing* in λ, meaning that in most states the value of becoming informed next period actually decreases as the current fraction informed increases, through the mechanisms discussed in Section 4.5. In the figure, this property holds wherever the value of becoming informed is positive — that is, wherever the curves are above the horizontal line at zero. (The same is true when we compare $J^I_t - J^U_t$ at different levels of λ.) Our model thus exhibits strategic substitutability in information acquisition wherever information acquisition is beneficial. Our model produces large price drops not because of strategic complementarity in information choices but rather despite an absence of strategic

[35] Here J^I_{t+1} is the value of being informed net of the future value of the cost of information Rc_I.

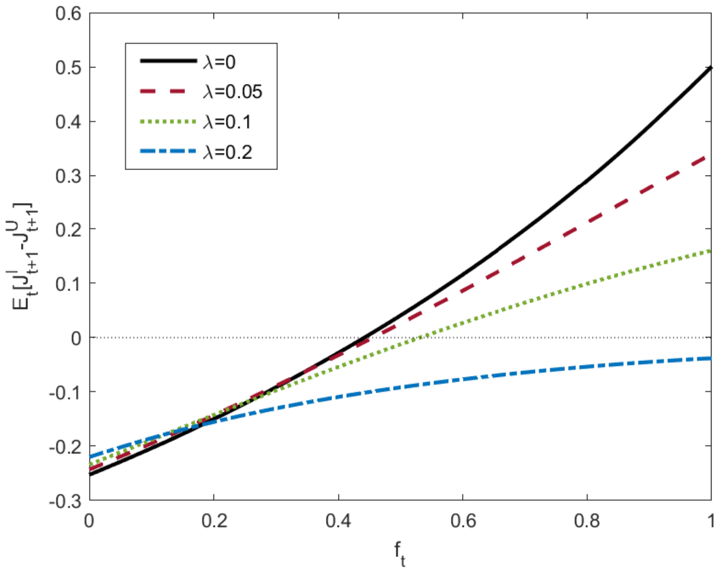

Figure 11: The curves show the value of information acquisition when λ is fixed exogenously. J_t^I is the value of being informed net of the future value of the cost of becoming informed Rc_I.

complementarity when λ is exogenous.[36]

We noted in Section 2.1 that the cost structure of information production may contribute to the feedback effect in the f_t dynamics (4). With high fixed costs for information production, an increase in demand for information can lower the price of information; holding the price of information acquisition constant, information providers may instead respond to increased demand by increasing information precision. The cost structure of information production is highlighted by Veldkamp (2006) as a source of strategic complementarity, so it is interesting to contrast the implications of our models. The models have different objectives and differ in many respects; but, most notably, in Veldkamp (2006) higher prices are associated with a larger fraction of informed investors, whereas we saw in Section 4.3 that in our model an increase in λ_t can precipitate a large price drop. The key difference is the dual role of information in our dynamic model discussed in Section 4.4: more precise information decreases dividend uncertainty but can increase future price variance. The second effect is absent in Veldkamp (2006), where investors earn dividends but do not earn capital gains from reselling their shares, making them indifferent to price variance. With no dependence on the next period's prices, the analysis reduces to a sequence of single-period problems.

To summarize, a static version of the feedback in the f_t dynamics (4), with $b_f\lambda(f_t)$ replaced

[36]This provides an alternative perspective on Figure 8, where we saw that fixing λ lowers the conditional return variance faced by uninformed investors, and decreases the price discount in the high information state.

with a constant $b_f\lambda$, does not produce complementarity in information acquisition. As discussed in Section 4.4–4.6, the price response to information generated by our model relies on the persistence of two information regimes and the dual role of information in a multiperiod model.

5.2 Changes in the Public-Private Mix ϕ_t

The model in (5) captures a different type of feedback mechanism. Recall that $1-\phi_t$ measures the fraction of knowable information that is public and costless to acquire. With $b_\phi < 0$ in (5), an increase in the fraction of informed investors λ_t increases the level of public information, suggesting a decrease in the value of acquiring costly information. In other words, the structure of (5) is more suggestive of strategic substitutability than complementarity. We are particularly interested in using the model to examine the effect of a positive shock to ϕ, corresponding to a loss of public information. We investigate the incentive this creates for investors to acquire costly non-public information to replace lost public information. The details of this analysis are in Appendix C, which may be read separately from most of the rest of this paper. Here we provide a brief summary.

As discussed in the introduction, the loss of sell-side analyst coverage resulting from the closures or mergers of brokerage research departments provides a setting with a plausibly exogenous loss of public information. The effects of such closures have been studied empirically in Hong and Kacperczyk (2010), Kelly and Ljungvist (2012), Balakrishnan et al. (2014), Johnson and So (2017), and Chen, Kelly and Wu (2018, CKW), among others. In particular, CKW study the effects of closures of brokerage research departments on the trading environment of affected stocks (i.e., those formerly covered by the closed research department). Importantly, the brokerage closures are driven by business conditions at the brokerages themselves and are unrelated to the prospects of the covered stocks. By focusing only on stocks that had five or fewer covering analysts prior to the brokerage closure, CKW identify securities whose information environment is particularly impacted by the loss of coverage. They document the following effects: price information efficiency falls; hedge funds trade more aggressively in the stock around earnings; hedge funds experience better investment performance in the stock; sophisticated investors increase their information acquisition; conditional on a large hedge fund presence in the stock, the loss of coverage has a smaller effect on price efficiency. Our model reproduces and helps explain these patterns.

CKW interpret their findings as indicative of "a substitution effect between sophisticated investors and providers of public information in facilitating market efficiency." As far as we know, ours is the first theoretical analysis to analyze this substitution effect in a dynamic setting. We model exogenous research department closures as positive shocks to ϕ_{t+1} in (5). The CKW substitution effect arises endogenously in our model because some traders choose to acquire costly information when the quality of public information deteriorates.

With $b_\phi < 0$, the presence of informed traders improves the public signal available to all investors in the next period. This can happen, for example, if hedge fund investors publicly disseminate their research ideas about a given stock and thereby provide de facto research reports to the public.[37] We show in Appendix C that the equilibrium effect of this feedback is to stabilize prices following exogenous brokerage closures, and we contrast our results with those of a static version of the model.

6 Conclusions

We have developed a model of a financial market in a dynamic information environment. The model combines exogenous shocks to the level of potentially available information, an endogenous response by investors, and feedback from investor information choices to the information environment. The dynamic structure of the model leads to a dual role for information, in which greater information reduces uncertainty about the next dividend but may increase price variance.

We study two types of feedback through which more information becomes available as more investors choose to become informed. In the first mechanism, increased demand increases the amount of potentially knowable (but costly) information; in the second mechanism, increased demand increases the level of public information through greater disclosures or information leakage. In both cases, changes in the level of information are accompanied by changes in the degree of information asymmetry between informed and uninformed investors.

We show that the equilibrium dynamics under the first mechanism, calibrated to market data, are characterized by two regimes, one with high prices and low volatility, and one with low prices and high volatility. A transition from the first regime to the second is reminiscent of a financial crisis but with no change in fundamentals — the price drop is driven by the dynamics of information and an increase in information asymmetry.

We use the second mechanism to explain empirical results of the effect on stock prices of the loss of information in the form of analyst coverage. In particular, we examine the substitution

[37] The Ira Sohn conferences, and the frequent appearances of hedge fund managers on CNBC, are examples of this effect. We could additionally model ϕ_t as a mean-reverting process to proxy for the possibility that firms either modify their information disclosures or get covered by the news media in a different way following exogenous ϕ_t shocks, though we do not pursue this idea further in this paper.

effect documented by Chen et al. (2018), showing that hedge funds respond to the loss of public information with greater private information acquisition. We show that accounting for information dynamics and incorporating feedback effects meaningfully change the predicted response to a loss of coverage when compared to a static model.

A Market Equilibrium

A.1 Proof of Proposition 3.1 (Existence of a Market Equilibrium)

Investor Demands for the Risky Asset

We prove Proposition 3.1 by solving explicitly for the coefficients of the price in (10). To allow for arbitrary variance beliefs, we write the investor optimization problem (8) as

$$\hat{J}_t^\iota \equiv \max_q \mathsf{E}\left[\mathsf{E}[W_{t+1}|\mathcal{I}_t^\iota, f_{t+1}, \phi_{t+1}] - \frac{\gamma}{2}\hat{\mathsf{var}}(W_{t+1}|\mathcal{I}_t^\iota, f_{t+1}, \phi_{t+1})\Big|\, \mathcal{I}_t^\iota\right], \quad \iota \in \{I, U\}, \tag{22}$$

where, using (15),

$$\hat{\mathsf{var}}(W_{t+1}|\mathcal{I}_t, f_{t+1}, \phi_{t+1}) = q^2(1+d)^2\left[\mathsf{var}(m_t|\mathcal{I}_t) + (1-f_t)\sigma_M^2\right] + q^2 V_B(f_{t+1}, \phi_{t+1}). \tag{23}$$

If the variance belief V_B is self-consistent, then (23) yields the conditional variance, but (22) makes explicit investors' objectives with arbitrary variance beliefs.

We can write the terminal wealth in (6) as $W_{t+1} = RW_t + q\pi_{t+1}$. Recalling that W_t is known to time-t investors, we set $\hat{\mathsf{var}}(\pi_{t+1}|\mathcal{I}_t, f_{t+1}, \phi_{t+1}) = \hat{\mathsf{var}}(W_{t+1}|\mathcal{I}_t, f_{t+1}, \phi_{t+1})/q^2$ to get

$$\hat{\mathsf{var}}(\pi_{t+1}|\mathcal{I}_t, f_{t+1}, \phi_{t+1}) = (1+d)^2\left[\mathsf{var}(m_t|\mathcal{I}_t) + (1-f_t)\sigma_M^2\right] + V_B(f_{t+1}, \phi_{t+1}). \tag{24}$$

The first-order condition for optimality in (22) becomes

$$q_t^\iota = \frac{1}{\gamma}\frac{\mathsf{E}\left\{\mathsf{E}[\pi_{t+1}|\mathcal{I}_t^\iota, f_{t+1}, \phi_{t+1}]|\, \mathcal{I}_t^\iota\right\}}{\mathsf{E}\left\{\hat{\mathsf{var}}(\pi_{t+1}|\mathcal{I}_t^\iota, f_{t+1}, \phi_{t+1})|\, \mathcal{I}_t^\iota\right\}} \equiv \frac{1}{\gamma}\frac{q_N^\iota}{q_D^\iota}, \tag{25}$$

where q_N^ι is the conditional expectation of the net profit, and q_D^ι is the expectation of its conditional variance, given a price variance of V_B. Through (24), the conditional variances, reflecting the cash-flow component and the time-t price variance, take the form

$$\begin{aligned} q_D^I &= (1+d)^2(1-f_t)\sigma_M^2 + \mathsf{E}_t V_B(f_{t+1}, \phi_{t+1}), \\ q_D^U &= q_D^I + (1+d)^2\mathsf{var}(m_t|P_t, \theta_t). \end{aligned} \tag{26}$$

Evaluating the conditional mean in the numerator of (25) as in (13), the demands for time-t informed and uninformed agents become

$$
\begin{aligned}
q^I &= \frac{q_N^I}{\gamma q_D^I} = \frac{1}{\gamma q_D^I}\Big[(1+d)(\mu_D + \rho D_t + \theta_t + m_t) + \mathsf{E}_t a(f_{t+1}, \phi_{t+1}) - RP_t\Big], \\
q^U &= \frac{q_N^U}{\gamma q_D^U} = \frac{1}{\gamma q_D^U}\Big[(1+d)(\mu_D + \rho D_t + \theta_t + \mathsf{E}[m_t|P_t, D_t, \theta_t]) + \mathsf{E}_t a(f_{t+1}, \phi_{t+1}) - RP_t\Big].
\end{aligned} \tag{27}
$$

For the informed, we have used the fact that $\mathsf{E}[m_t|\mathcal{I}_t^I] = m_t$ and $\mathsf{var}(m_t|\mathcal{I}_t^I) = 0$. For the uninformed, we evaluate (27) using[38]

$$
\begin{aligned}
\mathsf{E}[m_t|P_t, \theta_t] &= K_t(b_t m_t - c_t X_t), \\
\mathsf{var}(m_t|P_t, \theta_t) &= \phi_t f_t \sigma_M^2 (1 - K_t b_t) = \phi_t f_t \sigma_M^2 (1 - \mathcal{R}_t^2),
\end{aligned} \tag{28}
$$

with

$$
K_t = \frac{\mathsf{cov}(m_t, P_t|\theta_t, D_t)}{\mathsf{var}(P_t|\theta_t, D_t)} = \frac{b_t \phi_t f_t \sigma_M^2}{b_t^2 \phi_t f_t \sigma_M^2 + c_t^2 \sigma_X^2} \qquad \text{and} \qquad \mathcal{R}_t^2 \equiv K_t b_t. \tag{29}
$$

Market Clearing and Price Coefficients

We now impose market clearing (9), taking λ as given. We substitute investor demands q^ι in (9), use the price function from (10), and collect terms. We do not have to expand q_D^I or q_D^U in the following because these depend on (f_t, ϕ_t) but not on D_t, m_t, θ_t, or X_t. Equation (9) becomes

$$
\begin{aligned}
&\lambda q_D^U\Big[(1+d)(\mu_D + \rho D_t + \theta_t + m_t) + \mathsf{E}_t a(f_{t+1}, \phi_{t+1}) - RP_t\Big] \\
&+(1-\lambda) q_D^I\Big[(1+d)(\mu_D + \rho D_t + \theta_t + \mathsf{E}[m_t|P_t, \theta_t]) + \mathsf{E}_t a(f_{t+1}, \phi_{t+1}) - RP_t\Big] \\
&\qquad = \gamma q_D^I q_D^U X_t + \gamma q_D^I q_D^U \bar{X}.
\end{aligned} \tag{30}
$$

Collecting the D_t terms and then the θ_t terms yields the constants

$$
d = \frac{\rho}{R-\rho} \qquad \text{and} \qquad g = \frac{1}{R-\rho}. \tag{31}
$$

[38] Say $E[m|P] = K(bm + cX)$ for $K = \mathsf{cov}(m, P)/\mathsf{var}(P)$ and $\mathsf{var}(m|P) \equiv \mathsf{var}(m - E[m|P])$. Since $m - E[m|P] = (1-Kb)m - KcX$ then $\mathsf{var}(m - E[m|P]) = (1-Kb)^2 \mathsf{var}(m) + K^2 c^2 \mathsf{var}(X)$. This equals $(1 - 2Kb + K^2 b^2)\mathsf{var}(m) + K^2 c^2 \mathsf{var}(X) = (1 - 2Kb)\mathsf{var}(m) + K^2(b^2 \mathsf{var}(m) + c^2 \mathsf{var}(X))$. Note that $K = b\mathsf{var}(m)/(b^2 \mathsf{var}(m) + c^2 \mathsf{var}(X))$ and therefore $K^2(b^2 \mathsf{var}(m) + c^2 \mathsf{var}(X)) = b^2 \mathsf{var}^2(m)/(b^2 \mathsf{var}(m) + c^2 \mathsf{var}(X)) = Kb\mathsf{var}(m)$. And therefore $\mathsf{var}(m|P) = (1 - Kb)\mathsf{var}(m)$.

Collecting the constant terms in (30) yields

$$a_t = \frac{1}{R}\left[(1+d)\mu_D - \frac{\gamma q_D^I q_D^U}{\lambda q_D^U + (1-\lambda)q_D^I}\bar{X} + \mathsf{E}_t a(f_{t+1}, \phi_{t+1})\right]. \tag{32}$$

The function a appears on both sides. Assuming for a moment that a solution $a_t = a(f_t, \phi_t)$ exists, we can proceed to solve for b and c because a plays no role in the inference the uninformed make from the price in (29). We return to solve (32) after solving for b and c.

Collecting the m_t terms in (30) we get

$$b_t = \frac{1+d}{R} \times \frac{\lambda q_D^U + (1-\lambda)q_D^I K_t b_t}{\lambda q_D^U + (1-\lambda)q_D^I}. \tag{33}$$

Collecting the X_t terms and simplifying — mainly dividing the resulting equation by (33) — we find

$$b_t/c_t = \frac{\lambda(1+d)}{\gamma q_D^I}, \quad \text{with } c_t = \frac{\gamma q_D^U}{R} \text{ if } \lambda = 0. \tag{34}$$

We can now combine these equations to solve for b and c through the following steps, each of which involves only known quantities on the right side:

$$q_D^I(f,\phi) = (1+d)^2(1-f)\sigma_M^2 + \mathsf{E}[V_B(f_{t+1},\phi_{t+1})|f_t = f, \phi_t = \phi], \tag{35}$$

$$r(f,\phi) = \lambda(f,\phi)(1+d)/(\gamma q_D^I(f,\phi)), \tag{36}$$

$$\mathcal{R}^2(f,\phi) = \frac{r^2(f,\phi)f\phi\sigma_M^2}{r^2(f,\phi)f\phi\sigma_M^2 + \sigma_X^2}, \tag{37}$$

$$q_D^U(f,\phi) = q_D^I(f,\phi) + (1+d)^2 f\phi\sigma_M^2(1-\mathcal{R}^2(f,\phi)), \tag{38}$$

$$b(f,\phi) = \frac{1+d}{R}\frac{\lambda(f,\phi)q_D^U(f,\phi) + (1-\lambda(f,\phi))q_D^I(f,\phi)\mathcal{R}^2(f,\phi)}{\lambda(f,\phi)q_D^U(f,\phi) + (1-\lambda(f,\phi))q_D^I(f,\phi)}, \tag{39}$$

$$c(f,\phi) = \begin{cases} b(f,\phi)/r(f,\phi), & \lambda(f,\phi) > 0; \\ \gamma q_D^U(f,\phi)/R, & \lambda(f,\phi) = 0. \end{cases} \tag{40}$$

Equation (35) restates the first line of (26); (36) is the ratio in (34); (37) rewrites the expression for $\mathcal{R}^2$ in (29); (38) follows from the second line of (26); (39) and (40) come from (33) and (34).

A.2 Solving for the $a()$ curve

We now return to (32). Using the price function from (10) and the net profit from (7), we see that

$$\begin{aligned}
\mathsf{E}_t[\pi_{t+1}] &= \mathsf{E}_t[D_{t+1} + P_{t+1} - RP_t] \\
&= \mu_D + \rho D_t + \mathsf{E}_t[a_{t+1}] + d(\mu_D + \rho D_t) - Ra_t - RdD_t \\
&= (1+d)\mu_D + \mathsf{E}_t[a_{t+1}] - Ra_t \\
&= \gamma \bar{X} \frac{q_D^I q_D^U}{\lambda q_D^U + (1-\lambda) q_D^I}, \qquad (41)
\end{aligned}$$

where the third step follows from the definition of d in (31) and the fourth step follows from a_t in (32). Using (41) we can rewrite a_t in (32) as

$$\begin{aligned}
a_t &= \frac{1}{R}\left[(1+d)\mu_D - \mathsf{E}_t[\pi_{t+1}] + \mathsf{E}_t a(f_{t+1}, \phi_{t+1})\right] \\
&= \frac{1}{R}\left[(1+d)\mu_D - \mathsf{E}_t[\pi_{t+1}] + \mathsf{E}_t\left\{\frac{1}{R}\left[(1+d)\mu_D - \mathsf{E}_{t+1}[\pi_{t+2}] + \mathsf{E}_{t+1} a(f_{t+2}, \phi_{t+2})\right]\right\}\right] \\
&= \cdots = \frac{(1+d)\mu_D}{R-1} - \frac{1}{R}\sum_{i=0}^{\infty} \frac{1}{R^i} \mathsf{E}_t[\pi_{t+1+i}] \qquad (42)
\end{aligned}$$

If the variance belief V_B is bounded above and bounded away from zero, then $|\mathsf{E}_t[\pi_{t+1+i}]|$ is bounded and the expression in (42) is well-defined and finite. The quantities in (35)–(40) and (41) are all functions solely of the information state (f, ϕ), so the conditional expectations in (32) and (42) are taken with respect to the evolution of the information state in (4) or (5), for given λ. Equation (42) shows a_t is equal to the present value of all future expected dividend payments minus a discount reflecting the expected present value of all future net profits.

A.3 Statement and Proof of Proposition 3.2 (Existence of a Rational Expectations Equilibrium)

In this section, we prove the existence of a rational expectations equilibrium by showing that the variance belief updating mapping has a fixed point. This demonstrates the existence of self-consistent variance beliefs given an exogenously specified $\lambda()$ curve.[39]

We now precisely state Proposition 3.2 for model (4) with fixed ϕ and varying f. With f_t restricted to a finite set $\mathcal{D}$, we represent any function of f_t as a vector of dimension $n = |\mathcal{D}|$. We

[39] Note that this exogenous $\lambda()$ curve need not result in equivalent utilities for informed and uninformed investors. We endogenize the $\lambda()$ curve in Section B.

suppose $\lambda()$ is fixed (not necessarily constant) with $0 \leq \lambda(f,\phi) \leq 1$ for all f and ϕ. Let $F(b,c)$ be the mapping that sends the initial coefficients (b,c) to updated coefficients (b',c') through (16) and (35)–(40). A fixed point refers to the coefficients b and c such that $(b,c) = F(b,c)$.

Assume there exists a scalar $\bar{c} > 0$ satisfying the following four polynomial conditions:

$$\gamma\sigma_X\left(2\bar{q} + (1+d)^2\sigma_M^2\phi\right) - \bar{c}R\sigma_X\left(4 - (1+d)^2\gamma^2\sigma_M^2\sigma_X^2\phi\right) \leq 0, \tag{43}$$

$$4\gamma R\sigma_X^2\bar{c}\bar{q} - 4R^2\sigma_X^2\bar{c}^2 + (1+d)^2\sigma_M^2\phi\left(1+\gamma R\sigma_X^2\bar{c}\right)^2 \leq 0, \tag{44}$$

$$\gamma^2\sigma_X^2\left(\bar{q} + (1+d)^2\sigma_M^2\phi\right) \leq 1, \tag{45}$$

$$\gamma\left((1+d)^2\sigma_M^2\eta + \bar{c}^2\sigma_X^2\right) - R\bar{c} \leq 0, \tag{46}$$

where $\bar{q}$ in (43), (44), and (45) is a quadratic function in $\bar{c}$,

$$\bar{q} = (1+d)^2\sigma_M^2\delta + \bar{c}^2\sigma_X^2, \tag{47}$$

and the constants η and δ in (46) and (47) are defined by

$$\begin{aligned} \eta &:= \max_{f\in\mathcal{D}}\left\{\frac{1}{R^2}\mathsf{E}\left[f_{t+1}|f_t = f\right] + 1 - (1-\phi)f\right\}, \\ \delta &:= \max_{f\in\mathcal{D}}\left\{\frac{1}{R^2}\mathsf{E}\left[f_{t+1}|f_t = f\right] + 1 - f\right\}. \end{aligned} \tag{48}$$

Our simplest conditions would require $\delta \leq \eta \leq 2$, which is natural for $f \in [0,1]$, $\phi \in [0,1]$ and $R > 1$. The precise statement of Proposition 3.2 is that the mapping F has a fixed point in $[0,\bar{b}]^n \times [0,\bar{c}]^n$ with $n = |\mathcal{D}|$, $\bar{c}$ satisfying (43)–(46), and

$$\bar{b} = \frac{1+d}{R}.$$

The existence of a $\bar{c}$ satisfying (43)–(46), and thus the existence of fixed point for mapping F, only depends on model parameters.

As a shortcut for checking that these conditions hold, we show in Appendix A.3.1 that if $R \in [1, 1.2]$ and condition (17) holds, then F has a fixed point in $[0,\bar{b}]^n \times [0,\bar{c}]^n$ with $\bar{b} = (1+d)/R$ and $\bar{c} = R/(2\gamma\sigma_X^2)$.

For model (5) with fixed f and varying ϕ, a similar argument ensures the existence of a fixed point. Redefine the quantities δ and η as

$$\eta := \frac{1}{R^2}f + 1 \text{ and } \delta := \frac{1}{R^2}f + 1 - f.$$

Then we claim a fixed point exists in $[0,\bar{b}]^n \times [0,\bar{c}]^n$ as long as there exists $\bar{c} > 0$ satisfying (43)–(46) with ϕ replaced by f in these equations. Furthermore, assuming $R \in [1, 1.2]$, we can show using arguments similar to those in Section A.3.1 that (17) is a still sufficient condition for the existence of a fixed point.

Proof of Proposition 3.2

We prove the proposition for model (4) with fixed ϕ and varying f, then the result for model (5) with fixed f and varying ϕ follows in similar way. To prove the result, we use Brouwer's fixed point theorem, which states that if F is a continuous function mapping a compact convex set S to itself, then F has a fixed point in this set, meaning there exists $x \in S$ for which $x = F(x)$; see, for example, p.29 of Border (1989).

We show conditions for the Brouwer's fixed point theorem are satisfied with the belief updating mapping F and the compact convex set $[0,\bar{b}]^n \times [0,\bar{c}]^n$, $n = |\mathcal{D}|$. First, it is evident that F is continuous as each mapping in (35)–(40) is continuous in its input. It is also evident that $F(b,c) \geq 0$ since each step in (35)–(40) returns a nonnegative value. Next, for any input (b,c), we have $b'(f,\phi) \leq (1+d)/R = \bar{b}$ for all f and ϕ: it is easy to see $q^I_D(f,\phi) \geq 0$, $\mathcal{R}^2(f,\phi) \in [0,1]$, and $q^U_D(f,\phi) \geq 0$, so the second factor in (39) is in $[0,1]$ and the bound on $b'(f,\phi)$ follows.

It only remains to show that if $(b,c) \in [0,\bar{b}]^n \times [0,\bar{c}]^n$ and $(b',c') = F(b,c)$, then $c'(f,\phi) \leq \bar{c}$ for all f and ϕ. For this we first establish two useful bounds. To lighten notation, in the following we abbreviate conditional expectations of the form $\mathsf{E}[V_B(f_{t+1},\phi)|f_t = f]$ as $\mathsf{E}_t[V_B(f,\phi)]$. Recalling (16), we get the bound

$$\begin{aligned}
\mathsf{E}_t[V_B(f,\phi)] &= \mathsf{E}_t\left[b(f,\phi)^2 f\right]\phi\sigma_M^2 + g^2(1-\phi)\mathsf{E}_t[f]\sigma_M^2 + \mathsf{E}_t\left[c(f,\phi)^2\right]\sigma_X^2 \\
&\leq \bar{b}^2\mathsf{E}_t[f]\phi\sigma_M^2 + g^2(1-\phi)\mathsf{E}_t[f]\sigma_M^2 + \bar{c}^2\sigma_X^2 \\
&= (1+d)^2\frac{1}{R^2}\mathsf{E}_t[f]\sigma_M^2 + \bar{c}^2\sigma_X^2. \qquad (49)
\end{aligned}$$

The inequality uses the assumption $(b,c) \in [0,\bar{b}]^n \times [0,\bar{c}]^n$, and the last equality uses the relationship $\bar{b} = (1+d)/R = g$, which follows (31). Combining this with (35), we can bound $q^I_D(f,\phi)$ by

$$\begin{aligned}
q^I_D(f,\phi) &= (1+d)^2(1-f)\sigma_M^2 + \mathsf{E}_t[V_B(f,\phi)] \\
&\leq (1+d)^2\sigma_M^2\delta + \bar{c}^2\sigma_X^2 = \bar{q}, \qquad (50)
\end{aligned}$$

where $\bar{q}$ is given in (47). The inequality follows from the definition of δ in (48). We now proceed to prove the desired bound $c'(f,\phi) \le \bar{c}$ for all f and ϕ by the following two cases.

I. The Case Without Informed Investor: $\lambda(f,\phi) = 0$

We first prove the bound for $c'(f,\phi)$ for the case $\lambda(f,\phi) = 0$. By the second case in (40), $c'(f,\phi)$ is now given by $c'(f,\phi) = \gamma q_D^U(f,\phi)/R$. With $\lambda(f,\phi) = 0$, (36), (37), and (39) lead to $r(f,\phi) = \mathcal{R}^2(f,\phi) = b'(f,\phi) = 0$. Plugging these into (35) and (38) we have

$$
\begin{aligned}
q_D^U(f,\phi) &= q_D^I(f,\phi) + (1+d)^2 f\phi\sigma_M^2 \\
&= (1+d)^2(1-f)\sigma_M^2 + \mathsf{E}_t[V_B(f,\phi)] + (1+d)^2 f\phi\sigma_M^2.
\end{aligned}
$$

With the bound for $\mathsf{E}_t[V_B(f,\phi)]$ in (49), we can derive

$$
\begin{aligned}
q_D^U(f,\phi) &\le (1+d)^2\sigma_M^2\left(1 - f + \frac{1}{R^2}\mathsf{E}_t[f] + f\phi\right) + \bar{c}^2\sigma_X^2 \\
&\le (1+d)^2\sigma_M^2\eta + \bar{c}^2\sigma_X^2,
\end{aligned}
$$

where the last inequality follows from the definition of η in (48). Then the updated coefficient $c'(f,\phi)$ satisfies

$$
c'(f,\phi) = \frac{\gamma}{R} q_D^U(f,\phi) \le \frac{\gamma}{R}\left((1+d)^2\sigma_M^2\eta + \bar{c}^2\sigma_X^2\right) \le \bar{c},
$$

which follows from condition (46).

II. The Case With Informed Investors: $\lambda(f,\phi) > 0$

Next, we prove the bound $c(f,\phi) \le \bar{c}$ under the case $\lambda(f,\phi) > 0$. Applying the first case in (40), we get

$$
c'(f,\phi) = \frac{b(f,\phi)}{r(f,\phi)} = \frac{\gamma}{R}\left(\frac{q_D^I(f,\phi)}{\lambda(f,\phi)} \cdot \frac{\lambda(f,\phi)q_D^U(f,\phi) + (1-\lambda(f,\phi))q_D^I(f,\phi)\mathcal{R}^2(f,\phi)}{\lambda(f,\phi)q_D^U(f,\phi) + (1-\lambda(f,\phi))q_D^I(f,\phi)}\right). \tag{51}
$$

We need to derive a bound for the quantity in the parenthesis above.

We substitute $q_D^U(f,\phi)$ and $\mathcal{R}^2(f,\phi)$ in (51) using (38), (36), and (37). Then the right-hand side of (51) can be expressed as

$$
\frac{\gamma}{R}\left(\frac{\lambda q_D^I(1+d)^2\sigma_M^2 f\phi + \left(q_D^I\right)^2\gamma^2\sigma_X^2(1+d)^2\sigma_M^2 f\phi + (q_D^I)^3\gamma^2\sigma_X^2}{\lambda^2(1+d)^2\sigma_M^2 f\phi + \lambda q_D^I\gamma^2\sigma_X^2(1+d)^2\sigma_M^2 f\phi + (q_D^I)^2\gamma^2\sigma_X^2}\right), \tag{52}
$$

where we have dropped the dependence on f and ϕ to simplify notation. In the following, we establish the needed bound for $c'(f,\phi)$ by combining expression (52), bound for $q_D^I(f,\phi)$ in (50), and conditions (43)–(45). In particular, we need a bound of (52) that is valid for all $\lambda \in (0,1]$, as we do not make λ endogenous in this proof.

We first consider the case $f=0$ or $\phi=0$. If this holds, (52) directly reduces to $\gamma q_D^I(f,\phi)/R$ itself. By (50), $c'(f,\phi)$ satisfies

$$c'(f,\phi) = \frac{\gamma}{R} q_D^I(f,\phi) \leq \frac{\gamma}{R}\bar{q} = \frac{\gamma}{R}\left((1+d)^2\sigma_M^2\delta + \bar{c}^2\sigma_X^2\right).$$

By (48), it is easy to verify $\delta \leq \eta$ always holds given $\phi \geq 0$. Thus the desired bound for $c'(f,\phi)$ can be established as

$$c'(f,\phi) \leq \frac{\gamma}{R}\left((1+d)^2\sigma_M^2\delta + \bar{c}^2\sigma_X^2\right) \leq \frac{\gamma}{R}\left((1+d)^2\sigma_M^2\eta + \bar{c}^2\sigma_X^2\right) \leq \bar{c},$$

where the last inequality directly follows from condition (46).

Next, we consider the case (52) for $f>0$ and $\phi>0$. Here we need to consider the maximum of (52) over all $\lambda \in (0,1]$. We compute the derivative of (52) with respect to λ. As the numerator and denominator of (52) are linear and quadratic functions in λ respectively, the numerator of its derivative is a quadratic function in λ. Furthermore, it is easy to check this quadratic function is concave and the denominator of the derivative is always positive. Thus the maximum of (52) is attained at the larger root of its derivative, as long as that root falls in $(0,1]$. Through algebraic manipulations, the larger root of the derivative is given by

$$\tilde{\lambda} = \frac{-\tau + q_D^I\gamma\sigma_X\sqrt{\tau + (1+d)^2\sigma_M^2 f\phi}}{(1+d)^2\sigma_M^2 f\phi}, \tag{53}$$

where $\tau = q_D^I\gamma^2\sigma_X^2\left(q_D^I + (1+d)^2\sigma_M^2 f\phi\right)$. Also, $\tilde{\lambda} \in (0,1]$ is guaranteed by condition (45) and the bounds $q_D^I \leq \bar{q}$ and $f \leq 1$. Thus the maximum of (52) is indeed attained at $\lambda = \tilde{\lambda}$.

Letting $\lambda = \tilde{\lambda}$ in (52), we can derive

$$c'(f,\phi) \leq \frac{\gamma\sigma_X\left(2q_D^I + (1+d)^2\sigma_M^2 f\phi\right) + 2\sqrt{(1+d)^2\sigma_M^2 f\phi + q_D^I\gamma^2\sigma_X^2\left(q_D^I + (1+d)^2\sigma_M^2 f\phi\right)}}{R\sigma_X(4-(1+d)^2\gamma^2\sigma_M^2\sigma_X^2 f\phi)}. \tag{54}$$

Denote the right side by $\kappa(f, q_D^I)$. It is positive as condition (43) directly implies the denominator is greater than zero. As $\kappa(f,q_D^I)$ clearly increases in both q_D^I and f, it can be further bounded by setting q_D^I and f at their upper bounds $\bar{q}$ and one, respectively. Thus to establish the needed bound $c'(f,\phi) \le \bar{c}$, it suffices to show

$$c'(f,\phi) \le \kappa(f, q_D^I) \le \kappa(1,\bar{q}) \le \bar{c}. \tag{55}$$

Setting $q_D^I = \bar{q}$ and $f = 1$ on the right side of (54), the needed inequality $\kappa(1,\bar{q}) \le \bar{c}$ becomes

$$\begin{aligned} 2\sqrt{(1+d)^2\sigma_M^2\phi + \bar{q}\gamma^2\sigma_X^2\left(\bar{q} + (1+d)^2\sigma_M^2\phi\right)} \quad \le \quad & \bar{c}R\sigma_X\left(4 - (1+d)^2\gamma^2\sigma_M^2\sigma_X^2\phi\right) \\ & -\gamma\sigma_X\left(2\bar{q} + (1+d)^2\sigma_M^2\phi\right). \end{aligned} \tag{56}$$

By condition (43), the right side of (56) is positive, so this inequality is equivalent to the one obtained by squaring both sides. Taking squares and simplifying, we can show the quadratic terms $\bar{q}^2$ cancel out, and inequality (56) is equivalent to

$$4\gamma R\sigma_X^2\bar{c}\bar{q} \le 4R^2\sigma_X^2\bar{c}^2 - (1+d)^2\sigma_M^2\phi\left(1 + \gamma R\sigma_X^2\bar{c}\right)^2,$$

which holds as long as $\bar{c}$ and $\bar{q}$ satisfy condition (44). By (55), this proves the needed bound $c'(f,\phi) \le \bar{c}$ for $\lambda(f,\phi) > 0$. Combining the two cases with $\lambda(f,\phi) = 0$ and $\lambda(f,\phi) > 0$, we have proved $c' \le \bar{c}$ holds when conditions (43)–(46) are satisfied. The existence of a fixed point for the variance belief updating mapping now follows by Brouwer's theorem.

A.3.1 The Simplified Condition in (17)

Finally, we prove that (17) is a simple sufficient condition for the existence of fixed point. We follow the general conclusions established above and show that as long as $R \in [1, 1.2]$ and condition (17) hold, then the value $\bar{c} = R/(2\gamma\sigma_X^2)$ satisfies conditions (43)–(46). Thus by the proposition, a fixed point of belief updating exists in $[0, (1+d)/R]^n \times [0, R/(2\gamma\sigma_X^2)]^n$.

Plugging $\bar{c} = R/(2\gamma\sigma_X^2)$ into (43)–(46), algebraic computation shows that they simplify to following equivalent conditions:

$$(1+d)^2\gamma^2\sigma_M^2\sigma_X^2 \le \frac{3R^2}{3R^2 + 4\delta + 6}, \quad (1+d)^2\gamma^2\sigma_M^2\sigma_X^2 \le \frac{2R^4}{(4 + 4R^2 + R^4)\phi + 8R^2\delta}, \tag{57}$$

$$(1+d)^2\gamma^2\sigma_M^2\sigma_X^2 \le \frac{4 - R^2}{4(\delta + \phi)}, \text{ and } (1+d)^2\gamma^2\sigma_M^2\sigma_X^2 \le \frac{R^2}{4\eta}, \tag{58}$$

respectively. A nice property of these new conditions is that they are all upper bounds for the product $(1+d)\gamma\sigma_M\sigma_X$. Thus it suffices to propose a bound for $(1+d)\gamma\sigma_M\sigma_X$ that is small enough such that all the four conditions are satisfied. Since all the bounds in the right-hand sides of (57) and (58) are decreasing in ϕ, δ, and η, it suffices to consider their values at $\phi = 1$

and $\delta = \eta = 2$, as by (48) we clearly have $\delta \leq \eta \leq 2$. On the other hand, as the bounds do not monotonically depend on the risk-free return R, we impose a relatively mild condition $R \in [1, 1.2]$. Plugging these values into the right-hand sides of (57) and (58), the minimum of the four upper bounds approximately equals to 0.283, thus (17) is sufficient for the four conditions. Consequently, the corresponding $\bar{c} = R/(2\gamma\sigma_X^2)$ is a valid upper bound for the existence of fixed point.

B Information Equilibrium

B.1 Equating Expected Utilities of Informed and Uninformed given V_B

In this section we discuss the procedure for solving for an endogenous λ given variance beliefs. Given the demands in (26) and variance beliefs, we have, for $\iota \in \{I, U\}$,

$$\begin{aligned} \mathsf{E}\big[\mathsf{E}[W_{t+1}|\mathcal{I}_t^\iota, f_{t+1}, \phi_{t+1}]\big|\mathcal{I}_t^\iota\big] &= q^\iota \times q_N^\iota + RW_t = \frac{(q_N^\iota)^2}{\gamma q_D^\iota} + RW_t, \\ \mathsf{E}\big[\mathsf{var}(W_{t+1}|\mathcal{I}_t^\iota, f_{t+1}, \phi_{t+1})\big|\mathcal{I}_t^\iota\big] &= (q^\iota)^2 \times q_D^\iota = \frac{(q_N^\iota)^2}{\gamma^2 q_D^\iota}. \end{aligned} \tag{59}$$

We can therefore write the agent's value function in (8), conditional on $\mathcal{I}_t^\iota$ as

$$J_t^\iota = RW_t + \frac{1}{2\gamma}\frac{(q_N^\iota)^2}{q_D^\iota}, \quad \iota \in \{I, U\}. \tag{60}$$

To find an endogenous λ in the sense of Definition 2.1, we need to evaluate the conditional expectation of $J_t^I - J_t^U$ given the information state (f_t, ϕ_t). We can pull the denominators q_D^ι out of the conditional expectation because we know from (35) and (38) that they are purely functions of the information state. For the numerator terms, using the demands from (27), the price process from (10) and the condition on a_t in (32), it is straightforward to show that

$$q_N^\iota = \mathsf{E}_t[\pi_{t+1}] + (1+d)\mathsf{E}[m_t|\mathcal{I}_t^\iota] - Rb_t m_t + Rc_t X_t$$

where $\mathsf{E}_t[\pi_{t+1}]$ — which is a function of λ — is given by (41). The θ_t and D_t terms drop out, as do the terms involving a_t.[40] Note that q_N^ι equals the expected net profit $\mathsf{E}_t[\pi_{t+1}]$, which only conditions on f_t and ϕ_t, adjusted for the information set of agent ι.

[40] In particular, we do not need to evaluate $a(\cdot)$ to find the endogenous $\lambda(\cdot)$, which is useful in solving the model numerically.

Since $\mathsf{E}[m_t|\mathcal{I}_t^I] = m_t$ we have

$$\mathsf{E}[(q_N^I)^2|f_t,\phi_t] = (\mathsf{E}_t[\pi_{t+1}])^2 + (1+d-Rb_t)^2\phi_t f_t\sigma_M^2 + R^2c_t^2\sigma_X^2. \tag{61}$$

And from (28) we have $\mathsf{E}[m_t|\mathcal{I}_t^U] = K_tb_tm_t - K_tc_tX_t$. From this we have that

$$\begin{aligned}\mathsf{E}[(q_N^U)^2|f_t,\phi_t] &= (\mathsf{E}_t[\pi_{t+1}])^2 + [(1+d)K_tb_t - Rb_t]^2\phi_t f_t\sigma_M^2 + [Rc_t - (1+d)K_tc_t]^2\sigma_X^2 \\ &= (\mathsf{E}_t[\pi_{t+1}])^2 + [(1+d)K_t - R]^2(b_t^2\phi_t f_t\sigma_M^2 + c_t^2\sigma_X^2).\end{aligned} \tag{62}$$

Combining these expressions with J_t^ι in (60), we get an expression for the difference in conditional expectations

$$\Delta_{(f,\phi)} = \mathsf{E}[J_t^I - Rc_I|f_t = f, \phi_t = \phi] - \mathsf{E}[J_t^U|f_t = f, \phi_t = \phi]. \tag{63}$$

When this difference is positive, the marginal investor has an incentive to become informed. For a given $\{f,\phi\}$ we numerically solve for the $\lambda \in [0,1]$ which sets $\Delta_{f,\phi} = 0$. If this difference is always strictly positive we set $\lambda = 1$, and if it is always strictly negative we set $\lambda = 0$.

B.2 Proof of Proposition 3.3 (Existence of $\lambda()$ given Variance Beliefs)

To simplify notation, we focus on the case of dynamic f_t, as in (4) with $\phi_t \equiv \phi_0$ fixed. The argument for varying ϕ_t with fixed f_t works the same way. We therefore write Δ_f in (63) and omit ϕ_t from the conditioning information on the right.

Recall that we have restricted f_t to a finite set $\mathcal{D}$. We need to be more explicit about the mapping to $\mathcal{D}$ in (4). Suppose $\mathcal{D} = \{s_1,\ldots,s_n\} \subset [0,1]$. Partition the extended real line using

$$-\infty = c_0 < c_1 < \cdots < c_n < c_{n+1} = \infty,$$

and let $\Pi_{\mathcal{D}} : (c_j, c_{j+1}] \mapsto s_{j+1}$, $j = 0,1,\ldots,n$. We prove the proposition for this choice of $\Pi_{\mathcal{D}}$.

We can write the difference in expected utilities (63) as

$$\Delta_f = \frac{1}{2\gamma}\mathsf{E}\left[\frac{q_N^{I2}}{q_D^I} - \frac{q_N^{U2}}{q_D^U}\middle| f_t = f\right] - Rc_I = \frac{1}{2\gamma}\left(\frac{\mathsf{E}[q_N^{I2}|f_t = f]}{q_D^I(f)} - \frac{\mathsf{E}[q_N^{U2}|f_t = f]}{q_D^U(f)}\right) - Rc_I. \tag{64}$$

The terms on the right depend on the mapping $\lambda : \mathcal{D} \mapsto [0,1]$. However, for each f, Δ_f depends on λ only through $\lambda(f)$. This follows from the expressions in (35)–(40). We may therefore write Δ_f as $\Delta_f(\ell)$, with the interpretation that ℓ is the value of $\lambda(f)$. The proposition will follow once we show that $\Delta_f(\cdot)$ is continuous: given continuity, either $\Delta_f(\ell^*) = 0$ at some $\ell^* \in [0,1]$

(in which case we set $\lambda(f) = \ell^*$), or $\Delta_f(\ell) < 0$ for all $\ell \in [0,1]$ (in which case we set $\lambda(f) = 0$), or $\Delta_f(\ell) > 0$ for all $\ell \in [0,1]$ (in which case we set $\lambda(f) = 1$). This specification satisfies the conditions in Definition 2.1.

To establish continuity of Δ, we use the representation in (64). It is evident that, holding f fixed, each of the operations in (36)–(39) is continuous in $\ell = \lambda(f)$. But λ is also implicit in (35) through the conditional expectation of the variance belief, which takes the form

$$\mathsf{E}[V_B(f_{t+1})|f_t = s_i] = \sum_{s_j \in \mathcal{D}} \mathsf{P}(f_{t+1} = s_j | f_t = s_i) V_B(s_j).$$

With $\lambda(s_i) = \ell$, the transition probabilities take the form

$$\begin{aligned} \mathsf{P}(f_{t+1} = s_j | f_t = s_i) &= \mathsf{P}(a_f + b_f \ell + \kappa_f (f_t - a_f) + \epsilon_{f,t+1} \in (c_j, c_{j+1}] | f_t = s_i) \\ &= \mathsf{P}(\epsilon_{f,t+1} \in (c_j - [a_f + b_f \ell + \kappa_f(s_i - a_f)], c_{j+1} - [a_f + b_f \ell + \kappa_f(s_i - a_f)]]), \end{aligned}$$

which is the integral of the density of $\epsilon_{f,t+1}$ over the indicated interval and is therefore continuous in the endpoints and in ℓ. It follows that $\mathsf{E}[V_B(f_{t+1})|f_t = f]$ is continuous in $\lambda(f)$, and therefore that the mapping (35)–(39) is continuous in $\lambda(f)$, including, in particular, $q_D^I(f)$ and $q_D^U(f)$.

Next we turn to (40) and verify that $c(f)$ is continuous at $\lambda(f) = 0$. Using (51), we can write, for $\lambda(f) > 0$,

$$c(f) = \frac{\gamma}{R} q_D^I(f) \left(\frac{\lambda (1+d)^2 \sigma_M^2 f \phi + q_D^I(f) \gamma^2 \sigma_X^2 (1+d)^2 \sigma_M^2 f \phi + (q_D^I(f))^2 \gamma^2 \sigma_X^2}{\lambda^2 (1+d)^2 \sigma_M^2 f \phi + \lambda q_D^I(f) \gamma^2 \sigma_X^2 (1+d)^2 \sigma_M^2 f \phi + (q_D^I(f))^2 \gamma^2 \sigma_X^2} \right).$$

As $\lambda(f) \to 0$, we have $\mathcal{R}^2(f) \to 0$ and

$$c(f) \to \frac{\gamma}{R} \left((1+d)^2 \sigma_M^2 f \phi + q_D^I(f) \right) = \frac{\gamma}{R} q_D^U(f),$$

which coincides with the value specified for $c(f)$ in (40) at $\lambda(f) = 0$.

B.3 Proof of Proposition 3.4 (Existence of Information Equilibrium)

To prove the result, we apply Kakutani's fixed point theorem, which states the following (see, for example, p.72 of Border 1989). Let the domain S be a non-empty, compact and convex set, and let $F : S \mapsto 2^S$ be a set-valued function on S. Suppose $F(x)$ is non-empty and convex for all $x \in S$, and suppose that F has a closed graph (as defined shortly). Then F has a fixed point, meaning a point $x \in S$ for which $x \in F(x)$.

Let $G(b, c, \lambda)$ be the mapping that sends initial coefficients (b, c) to updated coefficients

(b', c') using λ through (35)–(40). Let

$$F(b, c, \lambda) = (b', c', \Lambda(b', c')) = (G(b, c, \lambda), \Lambda(G(b, c, \lambda)))$$

be the mapping that returns the updated (b', c') and the set of λs in $\Lambda(b', c')$. This representation is consistent with our algorithm: we first update (b, c) given λ; we then solve for the endogenous λ given the new (b', c'). The mapping F returns all of $\Lambda(b', c')$, rather than a single element.

Based on the proof of Proposition 3.2, we can restrict (b, c) to a domain $[0, \bar{b}]^n \times [0, \bar{c}]^n$, $n = |\mathcal{D}|$, in the sense that $G(\cdot, \cdot, \lambda)$ maps this set into itself for any λ. We may therefore take the domain of F to be $S = [0, \bar{b}]^n \times [0, \bar{c}]^n \times [0, 1]^n$, which is compact and convex. Moreover, for any $(b, c, \lambda) \in S$, $F(b, c, \lambda)$ is non-empty (by Propositions 3.1 and 3.3), and it is convex by the definition of $\Lambda(b, c)$.

It only remains to show that F has a closed graph. The closed graph property states that for any sequences $x_n \to x$, $y_n \to y$ with $y_n \in F(x_n)$ we have $y \in F(x)$. Because G is single-valued and continuous (as in the proof of Proposition 3.2), it suffices to show the following:

$$\text{if } \lambda' \notin \Lambda(b', c'), \text{ then } \lambda \notin \Lambda(b, c) \text{ for all } (b, c, \lambda) \text{ in a neighborhood of } (b', c', \lambda'). \tag{65}$$

We detail the case of model (4). Recall from the discussion surrounding (63) that when we solve for a $\lambda \in \Lambda_o(b, c)$, we may solve for each $\lambda(f)$, $f \in \mathcal{D}$, separately; the conditions on $\lambda(f)$ for different values of f do not interact. For each f, we look for a point at which

$$\Delta_f(\ell) \equiv \Delta_{b,c,f}(\ell) = \mathsf{E}[J_t^I - Rc | f_t = f] - \mathsf{E}[J_t^U | f_t = f]$$

crosses zero and set $\lambda(f) = \ell$; if zero is never crossed, we get a boundary case of $\lambda(f) = 0$ or 1. We have written $\Delta_{b,c,f}$ to emphasize that the utilities on the right are evaluated using (b, c).

We know from Section 3.3 that $\Delta_{b,c,f}(\ell)$ is continuous in ℓ for each f, and continuity in (b, c) follows similarly from (63). If $\Delta_{b,c,f}(\cdot)$ crosses zero, then we may define the first and last zero crossings by

$$\ell_{\min}(f) = \min\{\ell \in [0, 1] : \Delta_{b,c,f}(\ell) = 0\}, \quad \ell_{\max}(f) = \max\{\ell \in [0, 1] : \Delta_{b,c,f}(\ell) = 0\};$$

otherwise, set $\ell_{\min}(f) = \ell_{\max}(f) = 0$ if $\Delta_{b,c,f}(\ell) < 0$ for all ℓ, and $\ell_{\min}(f) = \ell_{\max}(f) = 1$ if $\Delta_{b,c,f}(\ell) > 0$ for all ℓ.

Returning to (65), it now follows that if $\lambda' \notin \Lambda(b', c')$ then it must be that $\lambda'(f) \notin [\ell_{\min}(f), \ell_{\max}(f)]$ for some f, again because the constraints on each $\lambda'(f)$ depend only on that

f. In particular, then, it must be that $\Delta_{b',c',f}(\ell) \neq 0$ for all $\ell \in [0, \lambda'(f)]$ or for all $\ell \in [\lambda'(f), 1]$; it suffices to consider the first case because a symmetric argument works for the second case. Suppose $\Delta_{b',c',f}(\lambda'(f)) < 0$; a symmetric argument applies if $\Delta_{b',c',f}(\lambda'(f)) > 0$. Then

$$\Delta_{b',c',f}(\ell) < 0 \text{ for all } \ell \in [0, \lambda'(f)]. \tag{66}$$

We claim that this holds in a neighborhood of (b', c', λ'). To argue by contradiction, suppose not. In other words, suppose that in any ϵ neighborhood of (b', c', λ') we can find a point $(b_\epsilon, c_\epsilon, \lambda_\epsilon)$ and an $\ell_\epsilon \in [0, \lambda_\epsilon(f)]$ with $\Delta_{b_\epsilon,c_\epsilon,f}(\ell_\epsilon) = 0$. Taking a sequence of ϵs decreasing to zero, gives us a sequence of such $(b_\epsilon, c_\epsilon, \lambda_\epsilon)$ and ℓ_ϵ, with $(b_\epsilon, c_\epsilon, \lambda_\epsilon) \to (b', c', \lambda')$. As the ℓ_ϵ take values in the compact set $[0, 1]$, they have a convergent subsequence. So, by taking a subsequence ϵ' of the original ϵ values, we get, for some ℓ_0, $(b_{\epsilon'}, c_{\epsilon'}, \lambda_{\epsilon'}, \ell_{\epsilon'}) \to (b', c', \lambda', \ell_0)$. And since $\ell_{\epsilon'} \leq \lambda_{\epsilon'}(f)$ for all ϵ', we have $\ell_0 \leq \lambda'(f)$. By the continuity of Δ,

$$\Delta_{b',c',f}(\ell_0) = \lim_{\epsilon' \to 0} \Delta_{b_{\epsilon'},c_{\epsilon'},f}(\ell_{\epsilon'}) = \lim_{\epsilon' \to 0} 0 = 0,$$

which contradicts (66). We have thus shown that (66) holds in a neighborhood of (b', c', λ'). But then $\lambda \notin \Lambda(b, c)$, for all (b, c, λ) in a neighborhood of (b', c', λ'), which is what we needed to show to prove the closed graph property.

C Changes in the Public-Private Mix ϕ_t

We apply the model in (5) to the consequences of brokerage closures, as described in Section 5.2.

C.1 Model Calibration

In calibrating (5), we use several of the parameters in Table 1. For the dividend process, we use the same values of ρ and σ_M as before, and we continue to set $\bar{D} = \bar{X} = 1$. Because we are modeling trading in a stock during a non-stressed market environment, the average turnover is lower. In the Supplementary Appendix, we show that the mean monthly turnover of the S&P 600 Small Cap index is 0.085 per month. Using (19), we get $\sigma_X = 0.0533$. We use the Small Cap index because the loss of an analyst is less significant for large companies.

With R and γ as before, the annualized excess return, $12 \times \mathsf{E}_t[\pi_{t+1}]/P_0$, is roughly 9.5%. This level is plausible for small stocks, which combine their market beta with a small-minus-big factor exposure. We set the per month cost of being informed at $c_I = 0.25$; this cost parameter pins down the equilibrium number of informed investors.

$\bar{D}$	$\bar{X}$	R	σ_X	ρ	σ_M	γ	c_I	f_0	$\mathsf{P}[\epsilon_\phi \neq 0]$	ϵ_ϕ	b_ϕ
1	1	1.0015	0.0533	0.967	0.0470	0.45	0.25	0.50	0.003×2	0.25	{-0.04,0}

Table 2: Calibrated parameters for model (5)

With ϕ_t changing over time, we fix $f_t \equiv f_0$. We use $f_0 = 0.5$, which is higher than the annual predictability of firm-level earnings (of 0.2) and of overall market earnings (of 0.3), but is lower than the annual predictability of S&P 500 dividends (of 0.64) (see Glasserman and Mamaysky 2018).

Finally, in the ϕ dynamics in (5), we use a three-point distribution taking values $\{-0.25, 0, 0.25\}$ with probabilities $\{0.003, 0.994, 0.003\}$. A shock of $\epsilon_\phi = 0.25$ is meant to represent the closure of a single brokerage house for firms that are covered by four brokers, a case consistent with CKW (who look at five brokers or less). We assume a change in coverage (either a closure or research coverage by a new brokerage) occurs with a 0.3% probability per month.[41] We first set $b_\phi = 0$ to focus only on the substitution-through-price channel. We then set $b_\phi = -0.04$ to explore the potential direct information provision role of informed speculators. Table 2 summarizes our model parameters.

Figure 12 summarizes our model equilibrium quantities as functions of the information state ϕ. The fraction informed $\lambda(\phi)$ ranges from zero to six percent. With $\phi_t = 0$, we have no informed investors (since all of $\tilde{m}_t$ is public), and the stock would need to experience two positive shocks to ϕ to enter the region where λ becomes positive. The b coefficient of the price (10) increases in ϕ since informed trade more aggressively on their signal when public information is worse. The c coefficient increases in ϕ because with higher information asymmetry between informed and uninformed supply shocks command a high risk premium. Finally, the unconditional risk premium is also increasing in ϕ (i.e., a is decreasing) and for the same reason. The scalloping in the a curve is due to the fact that with $b_\phi = 0$, conditional on a given starting point ϕ_0, all future ϕ_t's can only take on one of four values.[42]

[41] Say there is a 1% probability per year of a brokerage either closing or opening. Then per month the probability is $1 - 0.99^{1/12}$, assuming independence across months. If a typical firm has 5 covering brokers, and their closure outcomes are independent, then the probability that no firm closes in a month is $(1-(1-0.99^{1/12}))^4 \approx 1-0.003$. Similarly if there are four currently non-covering brokers with a 1% per year chance of initiating coverage, we get a 0.003 per month chance of coverage.

[42] For $b_\phi < 0$ no such effects are visible.

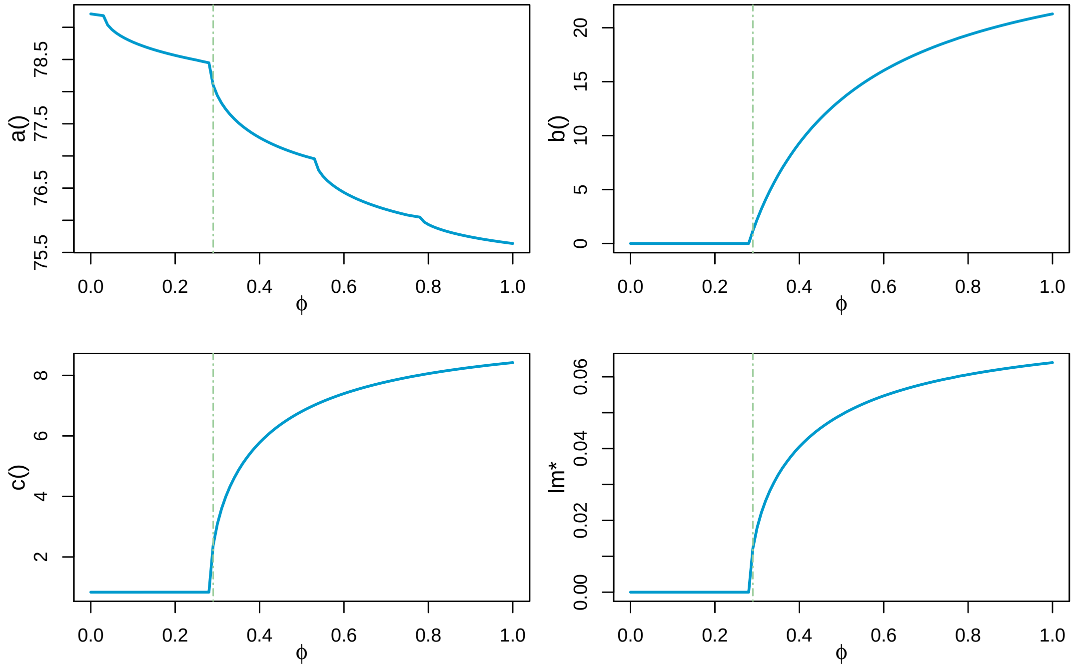

Figure 12: Equilibrium quantities for the ϕ_t model (5) using the parameters in Table 2. The vertical, dashed green line show the point at which $\lambda > 0$. Here $b_\phi = 0$.

$\Delta_V(\phi)$ for different model quantities V

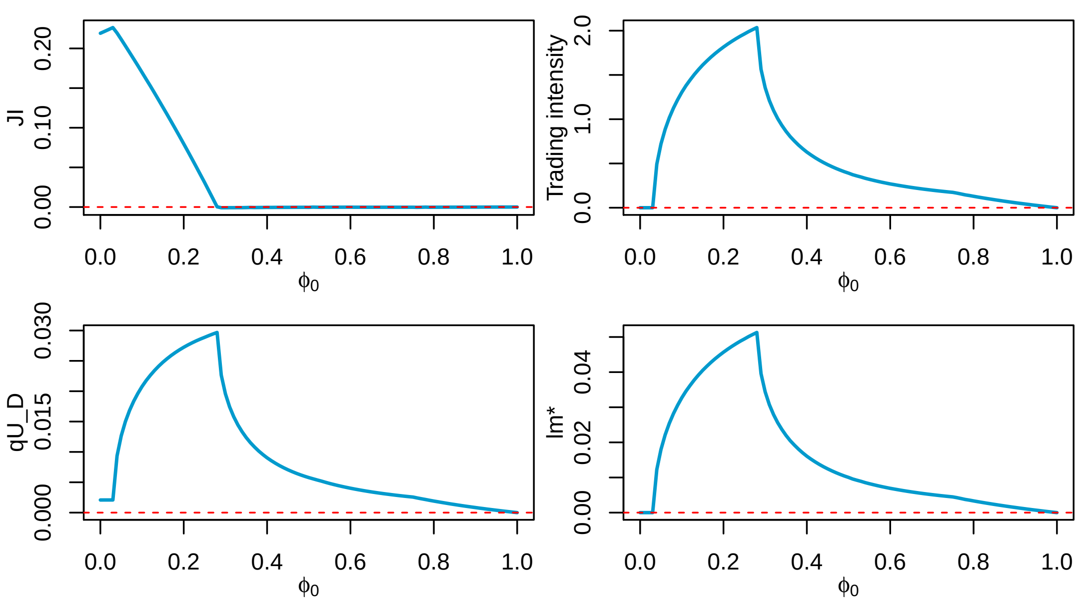

Figure 13: Equilibrium responses to a shock $\phi_{t+1} = \phi_t + 0.25$. The x-axis in the figures shows that value of ϕ_t immediately prior to the shock. Here $b_\phi = 0$.

C.2 Response to Brokerage Closures

We model a brokerage closure as a positive shock $\epsilon_\phi = 0.25$ to ϕ_{t+1}. For a given model variable V we then look at the difference in that variable as we jump from ϕ_t to ϕ_{t+1}, i.e.,

$$\Delta_V(\phi_t) \equiv V(\phi_t + \epsilon_\phi) - V(\phi_t).$$

Figure 13 shows the Δ_V's for the value function of the informed investors, for the model trading intensity (see below), for q_D^U (the expectation of the conditional variance of net profits for the uninformed, given by equation 25), and the equilibrium fraction informed λ. We now turn to CKW's five main results about stocks that experience a closure in a covering brokerage.

(1) CKW document that the magnitude of the post earnings announcement drift (PEAD) increases for affected stocks. They interpret PEAD as a proxy for the informational efficiency of the underlying stock. We interpret an increased PEAD in one of two ways. First, it indicates a higher price response from informed order flow, as measured by c_t in the model. We see that c_t increases with ϕ and therefore market liquidity $1/c_t$ (a measure suggested by Goldstein and Yang 2017, among others) falls. Second, the higher PEAD indicates a higher level of uncertainty faced by uninformed investors as to whether the current price is "fair." This is measured in our setting by the expected conditional variance of net profit π faced by the uninformed, q_D^U. From Figure 13, we see that $\Delta_{q_D^U}$ is always positive.

(2) CKW then show the hedge funds as a group trade more aggressively on stocks affected by brokerage closures around those stocks' earnings releases. In our case, we would like to measure how aggressively informed investors trade on their signal m_t. Here we appeal to the Goldstein and Yang (2015) *trading intensity* measure. They define trading intensity as

$$I_t = \lambda \frac{\partial q_t^I(m_t, P_t)}{\partial m_t},$$

which measures the sensitivity of all informed traders' demand (from equation 27) to changes in their signal m_t. In our model this trading intensity is given by[43]

$$I_t = \frac{b_t}{c_t}.$$

We see from Figure 13 that the trading intensity increases with reduced public information.

[43] Market clearing requires that $\lambda q^I(m,P) + (1-\lambda)q^U(P) = X + \bar{X}$. Differentiating with respect to X and using subscripts to indicate partial derivatives, we get $\lambda q_P^I P_X + (1-\lambda)q_P^U P_X = 1$ from which we conclude $P_X = 1/(\lambda q_P^I + (1-\lambda)q_P^U)$. Differentiating the market clearing condition with respect to m we get $\lambda(q_m^I + q_P^I P_m) + (1-\lambda)q_P^U P_m = 0$. This implies $\lambda q_m^I = -(\lambda q_P^I + (1-\lambda)q_P^U)P_m = -P_m/P_X = b/c$.

(3) Related to their increased trading intensity, CKW find that hedge funds experience better investment performance post closure. We have a direct measure of this in our model. From (59) and (60) we see that an investor's value function, J^{ι}_t, is equal to (scaled) expected wealth. Imagine the economy is in the no-informed region, but there is a small hedge fund whose cost of becoming informed is very low relative to other informed investors, and which therefore optimally participates in trading the stock. We assume further that the hedge fund trades with no price impact (to be consistent with being in the no-informed equilibrium). This small hedge fund's expected profits meaningfully increase with positive ϕ shocks, as can be seen from the top-left panel of Figure 13, as long as ϕ_t is in the no-informed region. Once the economy is in an interior equilibrium with $\lambda > 0$ the effect of a positive ϕ shock is to slightly lower the expected profits to being informed: Δ_{J^I} becomes negative, though it is indistinguishable from zero at the resolution of the figure. The decrease in J^I is due to a loss of risk sharing with the loss of public information.

(4) Furthermore, CKW find that sophisticated investors scale up their information acquisition following a brokerage research department closure. Following a positive ϕ shock we see that λ increases, as informed investors enter the market given the higher fraction of private information.

(5) Finally, CKW find that conditional on a large hedge fund presence in the stock, brokerage closure has a smaller effect on PEAD. As we see from Figure 12 there is a monotonic relationship between ϕ_t and λ_t. Therefore we can interpret the x-axis in Figure 13 as a proxy for the number of informed market participants. Conditional on having more informed participants, the impact of a ϕ shock on all equilibrium quantities we examine is decreasing.

Our model is therefore able to reproduce all five of the major findings in CKW. While our model is much more general than this particular example, its ability to capture a variety of empirical findings associated with an exogenous information shock is significant. We now examine more closely the implications of the dynamics of ϕ_t.

C.3 The Effect on Prices: The Stabilizing Role of Hedge Funds

Kelly and Ljungqvist (2012) study the short-term effects of a brokerage closure on the affected firms' stock prices. In an event study they find that affected firms experience cumulative abnormal returns (CAR) of −112 basis points on the day of an exogenous termination and of −261 basis points over the ensuing 5 days. We study the price response to a shock of $\epsilon_\phi = 0.25$ shock in our model under alternative ϕ_t dynamics.

To set a benchmark, we examine first a fully static model with $b_\phi = \epsilon_{\phi,t+1} = 0$. A ϕ shock

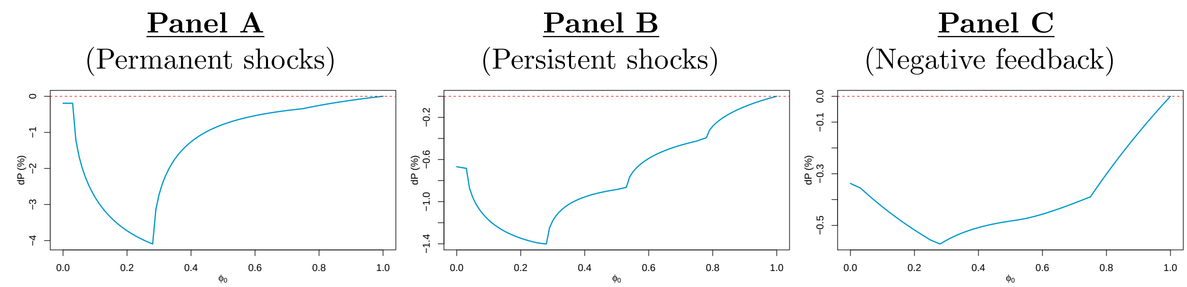

Figure 14: Equilibrium price responses to exogenous information shocks to ϕ. The x-axis in the charts shows the value of ϕ_t immediately prior to the shock. Panel A shows the case of a permanent change in ϕ_t with shocks ϵ_ϕ always at zero. Panel B shows the dynamic case with $b_\phi = 0$ and Panel C shows the dynamic case with $b_\phi = -0.04$.

is a one time, permanent and completely unanticipated change in the information environment. Figure 14 shows the percent drop in price at $\phi + \epsilon_\phi$ when the economy is at the steady-state level of all state variables. From the price function in (10), we see this is given by

$$100 \times \frac{a(\phi + \epsilon_\phi) - a(\phi)}{a(\phi) + d\bar{D}}.$$

Here the maximal next period (one-month in the calibration) price response is a price drop of 400 basis points (Panel A in Figure 14) at ϕ somewhere near 0.3.

We next examine the case from the prior section with $\mathsf{P}[\epsilon_{\phi,t+1} \neq 0] = 2 \times 0.003$ and $b_\phi = 0$. The maximal price response is now a drop of 140 basis points, as shown in Panel B. Therefore a static analysis meaningfully overstates the expected price drop due to exogenous information shocks, as accounting for the potential of offsetting future shocks meaningfully decreases the anticipated information asymmetry. Interestingly, the static ($b_\phi = \epsilon_{\phi,t+1} = 0$) and persistent ($b_\phi = 0$) cases, with maximal impacts of 400 and 140 basis points respectively, straddle the empirical finding in Kelly and Ljungqvist (2012) of a maximal impact of 261 basis points.

Finally, we allow for the feedback version of the model by setting $b_\phi = -0.04$. We interpret this as hedge funds sharing their research ideas with the general public either through investor presentations, appearance on news outlets, or unintended leakage of information. With $b_\phi = -0.04$, a 0.25 increase in ϕ from a starting point of $\phi = 0.25$ has a half-life along the equilibrium path of over six years (we iterate on equation 5 setting $\epsilon_{\phi,t+1} = 0$). So the information transmission from hedge funds to the public that we anticipate happens over very long periods of time. Yet even this very slow degree of reversion of ϕ back to its initial state meaningfully decreases the price impact of a ϕ shock, as can be seen in Panel C of Figure 14. The maximal percentage price drop is now only 50-60 basis points.

The intuition for this result is similar to the discussion of the price drop in Section ϕ_t can be interpreted as a measure of information asymmetry in the model conditional on the fixed level f of disclosure about the underlying dividend process. As before, agents are averse to high information asymmetry states — though information asymmetry in Section 4.3 was induced by an increase in f_t while here ϕ_t directly determines the asymmetry. In the static case of $b_\phi = \epsilon_{\phi,t+1} = 0$ agents believe the new higher level of asymmetry will last forever, and therefore the price concession demanded by the uninformed is very large. In the persistent case of $b_\phi = 0$ shocks to ϕ_t are persistent but can be reversed by future shocks in the opposite direction, and therefore are not permanent. In the feedback case of $b_\phi < 0$, shocks structurally diminish over time in response to an endogenous increase in informed. This meaningfully diminishes the anticipated information asymmetry, and therefore also the commensurate price impact.

Perhaps the large -261 basis point weekly CARs observed by Kelly and Ljungqvist (2012) reflects a lack of appreciation by the investment community of the stabilizing feedback effects from hedge fund participation that we posit here. Indeed our analysis suggests that hedge funds can play a stabilizing role in the case of exogenous shocks that increase information asymmetry by slowly disseminating their own research to the public.

References

Adrian, T. and Shin, H.S., 2010, "Liquidity and leverage," *J. Fin. Interm.* 19(3), 418–437.

Allen, F., Morris, S. and Shin, H.S., 2006, "Beauty contests and iterated expectations in asset markets," *Review of Financial Studies* 19(3), 719–752.

Admati, A., 1985, "A noisy rational expectations equilibrium for multi-asset securities markets," *Econometrica*, 53 (3), 629–657.

Avdis, E., 2016, "Information tradeoffs in dynamic financial markets," *Journal of Financial Economics*, 122, 568–584.

Balakrishnan, K., Billings, M.B., Kelly, B. and Ljungqvist, A., 2014, "Shaping liquidity: On the causal effects of voluntary disclosure," *Journal of Finance* 69(5), 2237–2278.

Barro, R.J., 2009, "Rare disasters, asset prices, and welfare costs," *American Economic Review* 99(1), 243–264.

Bernanke, B. and Gertler, M., 1989, "Agency costs, net worth, and business fluctuations," *American Economic Review* 79(1), 14–31

Biais, B., Bossaerts, P. and Spatt, C., 2010, "Equilibrium asset pricing and portfolio choice under asymmetric information," *Review of Financial Studies* 23(4), 1503–1543.

Bianchi, J., 2011, "Overborrowing and systemic externalities in the business cycle," *American Economic Review* 101(7), 3400–3426.

Border, K.C., 1989, *Fixed Point Theorems with Applications to Economics and Game Theory*, Cambridge University Press.

Boot, A., and A. Thakor, 2001, "The many faces of information disclosure," *Review of Financial Studies* 14(4), 1021–1057.

Botosan, C., 1997, "Disclosure level and the cost of equity capital," *Accntng. Rev.* 72, 323–349.

Brancati, E., and M. Macchiavelli, 2019, "The information sensitivity of debt in good and bad times," *Journal of Financial Economics* 133, 99–112.

Brennan, M.J. and Cao, H.H., 1997, "International portfolio investment flows," *Journal of Finance* 52(5), 1851–1880.

Campbell, J., 1991, "A variance decomposition for stock returns," *Economic J.* 101, 157–179.

Cesa-Bianchi, A. and E. Fernandez-Corugedo, 2018, "Uncertainty, financial frictions, and nominal rigidities: A quantitative investigation," *J. Money, Credit and Banking*, 50 (4), 603–636.

Chamley, C., 2007, "Complementarities in information acquisition with short-term trades," *Theoretical Economics* 2(4), 441–467.

Chen, S., Matsumoto, D. and Rajgopal, S., 2011, "Is silence golden? An empirical analysis of firms that stop giving quarterly earnings guidance," *Journal of Accounting and Economics* 51(1-2), 134–150.

Chen, Y., B. Kelly, and W. Wu, 2018, "Sophisticated investors and market efficiency: Evidence from a natural experiment," working paper.

Dang, T.V., Gorton, G. and Holmström, B., 2012, "Ignorance, debt and financial crises," working paper, Economics Department, Columbia University.

Dang, T.V., Gorton, G. and Holmström, B., 2019, "The information view of financial crises," *NBER Working Paper No. 26074*.

Dutta, S., and A. Nezlobin, 2017, "Information disclosure, firm growth, and the cost of capital," *Journal of Financial Economics* 123, 415–431.

Dow, J., I. Goldstein, and A. Guembel, 2017, "Incentives for information production in markets where prices affect real investment," *J. European Economic Association* 15(4), 877–909.

Fama, E. and K. French, 2000, "Forecasting profitability and earnings," *Journal of Business*, 73 (2), 161–175.

Frankel, R., S. P. Kothari, and J. Weber, 2006, "Determinants of the informativeness of analyst research," *Journal of Accounting and Economics* 41, 29–54.

Farboodi, M. and Veldkamp, L., 2017, "Long run growth of financial technology," Working Paper W23457, National Bureau of Economic Research.

Ganguli, J.V. and Yang, L., 2009, "Complementarities, multiplicity, and supply information," *Journal of the European Economic Association* 7(1), 90–115.

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

Glasserman, P., F. Li, and H. Mamaysky, 2019, "Time variation in the news-returns relationship," working paper, Columbia Business School.

Goldstein, I. and Leitner, Y., 2018, "Stress tests and information disclosure," *Journal of Economic Theory*, 177, 34–69.

Goldstein, I. and Yang, L., 2015, "Information diversity and complementarities in trading and information acquisition," *Journal of Finance*, 70, 1723–1765.

Goldstein, I. and Yang, L., 2017, "Information disclosure in financial markets," *Annual Review of Financial Economics* 9, 101–125.

Gorton, G. and Ordoñez, G., 2014, "Collateral crises," *Amer. Econ. Rev.* 104(2), 343–378.

Gorton, G. and Ordoñez, G., 2018, "Good booms, bad booms," *Journal of the European Economic Association*, to appear.

Grossman, S. and J. Stiglitz, 1980, "On the impossibility of informationally efficient markets," *American Economic Review*, 70(3), 393–408.

Healy, P.M., Hutton, A.P. and Palepu, K.G., 1999, "Stock performance and intermediation changes surrounding sustained increases in disclosure," *Contemp. Acc. Res.* 16(3), 485–520.

Hellwig, M., 1980, "On the aggregation of information in competitive markets," *Journal of Economic Theory*, 22, 477–498.

Hong, H. and Kacperczyk, M., 2010, "Competition and bias," *Quarterly Journal of Economics* 125(4), 1683–1725.

Johnson, T. and E. So, 2017, "A simple multimarket measure of information asymmetry," *Management Science*, 64 (3).

Kelly, B. and Ljungqvist, A., 2012, "Testing asymmetric-information asset pricing models," *Review of Financial Studies* 25(5), 1366–1413.

Kiyotaki, N. and Moore, J., 1997, "Credit cycles," *Journal of Political Economy* 105(2), 211–248

Kreps, D. and E. Porteus, 1978, "Temporal resolution of uncertainty and dynamic choice theory," *Econometrica*, 46 (1), 185–200.

Lo, A. and J. Wang, 2000, "Trading volume: Definitions, data analysis, and implications of portfolio theory," *Review of Financial Studies*, 13(2), 257–300.

Lorenzoni, G., 2008, "Inefficient credit booms," *Review of Economic Studies* 75(3), 809-833.

Manzano, C. and X. Vives, 2011, "Public and private learning from prices, strategic substitutability and complementarity, and equilibrium multiplicity," *Journal of Mathematical Economics* 47(3), 346–369.

Mondria, J., 2010, "Portfolio choice, attention allocation, and price comovement," *Journal of Economic Theory*, 145 (5), 1837–1864.

Ng, J., 2011, "The effect of information quality on liquidity risk," *Journal of Accounting and Economics* 52(2-3), 126–143.

Peress, J., 2010, "The tradeoff between risk sharing and information production in financial markets," *Journal of Economic Theory* 145(1), 124–155.

Rösch, D., A. Subrahmanyam, and M. A. van Dijk, 2017, "The dynamics of market efficiency," *Review of Financial Stuides* 30(4), 1151–1187.

Spiegel, M., 1998, "Stock price volatility in a multiple security overlapping generations model," *Review of Financial Studies* 11(2), 419–447.

Van Nieuwerburgh, S., and L. Veldkamp, 2010, "Information acquisition and portfolio underdiversification," *Review of Economic Studies* 77(2), 779–805.

Veldkamp, L., 2006, "Media frenzies in markets for financial information," *American Economic Review* 96(3), 577–601.

Vuolteenaho, T., 2002. "What drives firm-level stock returns?," *J. Finance* 57(1), 233–264.

Wang, J., 1993, "A model of intertemporal asset prices under asymmetric information," *Review of Economic Studies*, 60, 249–282.

Wang, J., 1994, "A model of competitive stock trading volume," *J. Pol. Econ.* 102, 127–168.

Watanabe, M., 2008, "Price volatility and investor behavior in an overlapping generations model with information asymmetry," *Journal of Finance* 63(1), 229–272.