

MIT Open Access Articles

Long-Run Growth of Financial Data Technology

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Farboodi, Maryam and Laura Veldkamp. "Long-Run Growth of Financial Data Technology." *American Economic Review*, 110, 8 (August 2020): 2485-2523 © 2020 The Author(s)

As Published: 10.1257/AER.20171349

Publisher: American Economic Association

Persistent URL: <https://hdl.handle.net/1721.1/129767>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Long-Run Growth of Financial Data Technology[†]

By MARYAM FARBOODI AND LAURA VELDKAMP*

“Big data” financial technology raises concerns about market inefficiency. A common concern is that the technology might induce traders to extract others’ information, rather than to produce information themselves. We allow agents to choose how much they learn about future asset values or about others’ demands, and we explore how improvements in data processing shape these information choices, trading strategies and market outcomes. Our main insight is that unbiased technological change can explain a market-wide shift in data collection and trading strategies. However, in the long run, as data processing technology becomes increasingly advanced, both types of data continue to be processed. Two competing forces keep the data economy in balance: data resolve investment risk, but future data create risk. The efficiency results that follow from these competing forces upend two pieces of common wisdom: our results offer a new take on what makes prices informative and whether trades typically deemed liquidity-providing actually make markets more resilient. (JEL C55, D83, G12, G14, O33)

In most sectors, technological progress boosts efficiency. But in finance, more efficient data processing and the new data-intensive trading strategies it spawned have been blamed for market volatility, illiquidity, and inefficiency. One reason financial technology is suspect is that its rise has been accompanied by a shift in the nature of financial analysis and trading. Instead of “kicking the tires” of a firm, investigating its business model or forecasting its profitability, many traders today engage in statistical arbitrage: they search for “dumb money” or mine order flow data and develop algorithms to profit from patterns in others’ trades. Why might investors choose one strategy over another and why are these incentives to process each type of data changing? Answering these questions requires a model. Just as past investment rates

*Farboodi: MIT Sloan and NBER (email: farboodi@mit.edu); Veldkamp: Columbia Graduate School of Business, NBER, and CEPR (email: lv2405@columbia.edu). John Leahy was the coeditor for this article. We thank three anonymous referees for their help and support. We also thank Daniel Andrei, Marios Angeletos, Markus Brunnermeier, Martin Eichenbaum, Pete Kyle, Sergio Rebelo, Steven Strongin, Robert Ulbricht, and Xavier Vives, seminar and conference participants at AFA, Barcelona GSE, Cornell, Fordham, Maryland, NYU, Princeton, Stanford, Yale, and the SED conference, the NASDAQ DRP research day, and the LAEF conference on information in finance for comments. We thank the Global Markets Institute at Goldman Sachs for their financial support through the GMI Fellowship program. Any opinions, findings, or conclusions are those of the authors and do not reflect the views of Goldman Sachs or the GMI. We thank John Barry, Chase Coleman, Matias Covarrubias, Adam Lee, Ran Liu, Roxana Mihet, and Arnav Sood for their capable research assistance.

[†]Go to <https://doi.org/10.1257/aer.20171349> to visit the article page for additional materials and author disclosure statements.

are unreliable forecasts for economies in transition, empirically extrapolating past financial trends is dubious in the midst of a technological transformation.

To make sense of current and future long-run trends requires a growth model of structural changes in the financial sector. Since much of the technological progress is in the realm of data processing, we use an information choice model to explore how unbiased technological progress changes what data investors choose to process, what investment strategies they adopt, and how the changing strategies alter financial market efficiency. Structural change in the financial sector arises because improvements in data processing trigger a shift in the type of data investors choose to process. Instead of data about firm fundamentals, firms choose to process more and more data about other investors' demand. Each data choice gives rise to an optimal trading strategy. The resulting shift in strategies is an abandonment of value investing and the rise of a strategy that is part statistical arbitrage, part retail market making, and part to extract what others know. Just like the shift from agriculture to industry, some of the data-processing shift we describe takes place because growing efficiency interacts with decreasing returns. But unlike physical production, information leaks out through equilibrium prices, producing externalities and a region of endogenous increasing returns that do not arise in standard growth models.

The consequences of this shift in strategy upend two pieces of common wisdom. First, the abandonment of fundamentals-based investing does not necessarily compromise financial market efficiency. Efficiency, as measured by price informativeness, continues to rise, even as fundamental data gathering falls. These results change the way we interpret empirical measurement. They support the interpretation of price informativeness as a proxy for total information, but not the idea that price informativeness measures information that is specifically about firm fundamentals. Our second surprise is that the price impact of an uninformed trade (liquidity) does not consistently fall. Even though demand data allow investors to identify uninformed trades, and even though investors use this information to "make markets" for demand-driven trades, market-wide liquidity may not improve.

A key theme of the paper and the force underlying the long-run results, including the stagnation of liquidity, is the following observation: data resolve risk, but they also create it. First, data resolve risk by allowing investors to better forecast asset payoffs. If more of the payoff is predictable, the remaining risk is less. That force is present in most information choice models. The opposing force arises because we use a dynamic model. Most models of information choice have assets that are one-period lived; at the start of the period, someone buys assets that offer an exogenous payoff at the end. In our model, as in equity markets, assets are multi-period lived. The payoff of an asset depends on its future price. But the future price depends on what the future investors will know. If future investors have a large amount of data, the future price is more informative and thus more sensitive to future innovations in dividends or in demand. If future prices are more sensitive to future news, they are riskier today. A future price that is less forecastable today creates risk, which we call *future information risk*. This competition between static risk-reduction and future information risk-creation governs the long-run market liquidity and information choices.

To think about choices related to information and their equilibrium effects, a noisy rational expectations framework is an obvious starting point. We add three

ingredients. First, we add a continuous choice between firm fundamental information and investor demand information. We model data processing in a way that draws on the information processing literatures in macroeconomics and finance.¹ However, the trade-off between processing data about trading demand and data about firm fundamentals is new to the literature. This trade-off is central to contemporary debate and essential to our main results. Second, we add long-run technological progress. It is straightforward to grow the feasible signal set. But doing so points this tool in a new direction, so that it answers a different set of questions. Third, we use long-lived assets, as in Wang (1993), because ignoring the fact that equity is a long-lived claim fundamentally changes our results. Because long-lived assets have future payoff shocks that are not learnable today, they ensure that uncertainty does not disappear and the equilibrium continues to exist, even when technology is advanced. Long-lived assets are also responsible for our long-run balanced growth path and the stagnation of liquidity. There are many aspects to the financial technology revolution and many details of modern trading strategies from which our analysis abstracts away. But building up a flexible equilibrium framework from well-understood building blocks, instead of a stylized, detail-oriented model, we give others the ability to adapt and use the framework to answer many long-run questions.

The key to the transition from a fundamental to a demand data strategy is understanding what makes each type of data valuable. Fundamental data are always valuable. They allow investors to predict future dividends and future prices. Demand data contain no information about any future cash flows. Yet, they have value because they enable an investor to trade against demand shocks, sometimes referred to as searching for “dumb money.” By buying when demand shocks are low and selling when they are high, an investor can systematically buy low, sell high, and profit. This is the sense in which the demand-data trading strategy looks like market making for uninformed retail investors. Demand-data processors stand ready to trade against, i.e., make markets for, uninformed orders. The mathematics of the model suggest another, complementary interpretation of the rise in demand-based strategies. Demand shocks are the noise in prices. Knowing something about this noise allows investors to remove that known component and to reduce the noise in prices. Since prices summarize what other investors know, removing price noise is a way of extracting others’ fundamental information. Seen in this light, the demand-driven trading strategy shares some of the hallmarks of automated trading strategies, largely based on order flow data, that are also designed to extract the information of other market participants.

Our main results in Section II describe the evolution of data processing in three phases.

Phase one: technology is poor and fundamental data processing dominates. In this phase, fundamental data are preferred because demand data have little value. To see why, suppose that no investor has any fundamental information. In such an environment, all trades are uninformed. No signals are necessary to distinguish between

¹See, e.g., Caplin, Leahy, and Matejka (2016); Mackowiak and Wiederholt (2012); Nimark (2008); Briggs et al. (2017); Kacperczyk, Nosal, and Stevens (2015); Barlevy and Veronesi (2000); Goldstein, Ozdenoren, and Yuan (2013); Blanchard, L’Huillier, and Lorenzoni (2013); Fajgelbaum, Schaal, and Taschereau-Dumouchel (2017); Angeletos and La’O (2013); Atkeson and Lucas (1992); Chien, Cole, and Lustig (2016); Bassetto and Galli (2017); Abis (2017).

informed and uninformed trades. As technology progresses and more trades are information-driven, identifying and trading against the remaining non-informational trades becomes more valuable.

Phase two: moderate technology generates increasing returns to demand data processing. Most physical production as well as most of the information choices in financial markets exhibit decreasing returns, also called strategic substitutability. Returns decrease because acquiring the same information as others leads one to buy the same assets that others do, and these assets are more expensive because they are popular. Our increasing returns come from an externality that is specific to data: information leaks through the equilibrium price. When more investors process demand data, they extract more fundamental information from equilibrium prices and they trade on that information. More trading on fundamental information, even if extracted, makes the price more informative, which encourages more demand data processing to enable more information extraction from the equilibrium price.

Phase three: high technology restores balanced data processing growth. As technology progresses, both types of data become more abundant. In the high-technology limit, both types grow in fixed proportion to each other. When information is abundant, the natural substitutability force in asset markets strengthens and overtakes complementarity.

The consequences of this shift in data analysis and trading strategies involve competing forces. Three economic forces (decreasing returns, increasing returns, and future information risk) appear whether technology is unbiased or biased, whether demand is persistent or not, and for most standard formulations of data constraints. We identify each force theoretically. However, if we want to learn which force is likely to dominate, we need to put some plausible numbers into the model. Section III calibrates the model to financial market data so that we can numerically explore the growth transition path and its consequences for market efficiency.

The market efficiency results call popular narratives into question. First, even as demand analysis crowds out fundamental analysis and reduces the discovery of information about the future asset value, price informativeness does not fall. The reason for this outcome is that demand information allows demand traders to extract fundamental information from prices. That makes the demand traders, and thus the average trader, better informed about future asset fundamentals. When the average trader is better informed, prices are more informative.

Second, even though demand traders systematically take the opposite side of uninformed trades, the rise of demand trading does not always enhance market liquidity. This is surprising because taking the opposite side of uninformed trades is often referred to as “providing liquidity.” This is one of the strongest arguments that proponents of activities such as high-frequency trading make to defend their methods. But if by providing liquidity, we really mean reducing the price impact of an uninformed trade, then the rise of demand trading may not accomplish that end. The problem is not demand trading today, but the expectation of future informed trading of any kind, fundamental or demand, creating future information risk. So future data processing raises the risk of investing in assets today. More risk per share of asset today is what causes the sale of one of the asset’s shares to have a larger effect on the price. Finally, the rise in demand-driven trading strategies, while it arises concurrent with worrying market trends, does not cause those trends. The rise in return

uncertainty, and the stagnation of liquidity, emerge as concurrent trends with financial data technology, not demand analysis, as their common cause.

Finally, Section IV explores suggestive evidence in support of the model. Section V concludes, offering ideas for future research.

Contribution to the Existing Literature.—Our model combines features from a few disparate literatures. Long-run trends in finance are featured in Asriyan and Vanasco (2014); Biais, Foucault, and Moinas (2015); Glode, Green, and Lowery (2012); and Lowery and Landvoigt (2016), which model growth in fundamental analysis or an increase in its speed. Caplin, Dean, and Leahy (2018) and Dávila and Parlatore (2016) explore a decline in trading costs. The idea of long-run growth in information processing is supported by the rise in price informativeness documented by Bai, Philippon, and Savov (2016).

A small, growing literature examines demand information in equilibrium models. In Ganguli and Yang (2009), agents can choose whether to purchase a fixed bundle of fundamental and demand information. In Yang and Zhu (2016) and Manzano and Vives (2010), the precision of fundamental and demand information is exogenous. Babus and Parlatore (2015) examines intermediaries who observe their customers' demands. Our demand signals also resemble Angeletos and La'O's (2013) sentiment signals about other firms' production; Banerjee and Green's (2015) signals about motives for trade; the signaling by He's (2009) intermediaries; and the noise in government's market interventions in Brunnermeier, Sockin, and Xiong (2017). But none of this research examines the choice central to this paper: whether to process more about asset payoffs or to analyze more demand data. Without that trade-off, previous studies cannot explore how trading strategies change as productivity improves. Furthermore, our research adds a long-lived asset to a style of model that has traditionally been static.² We do this because assets are not in fact static and assuming that they are reverses many of our results.

One interpretation of our demand information is that it is what high-frequency traders learn by observing order flow. Like high-frequency traders, our traders use data on asset demand to distinguish informed from uninformed trades, and they stand ready to trade against uninformed order flow. While our model has no frequency, making for a loose interpretation, it does contribute a perspective on this broad class of strategies. As such, it complements work by Du and Zhu (2017) and Crouzet, Dew-Becker, and Nathanson (2016) on the theory side, as well as empirical work such as Hendershott, Jones, and Menkveld (2011), which measures how fundamental and algorithmic trading affects liquidity. At the same time, if many high frequency trades are executed for the purpose of obscuring price information, our model does not capture this phenomenon. Such a practice could work in the opposite direction.

Another, more theoretical, interpretation of demand signals is that they make a public signal, the price, less conditionally correlated. Work by Myatt and Wallace (2012), Chahrour (2014), and Amador and Weill (2010) delves into similar choices between private, correlated, and public information that arise in strategic settings.

²Exceptions include 2- and 3-period models, such as Cespa and Vives (2012).

I. Model

To explore growth and structural change in the financial economy, we use a noisy rational expectations model with three key ingredients: a choice between fundamental and demand data, long-lived assets, and unbiased technological progress in data processing. A key question is how to model structural change in trading activity: the way in which investors earn profits has changed. A hallmark of that change is the rise in information extraction from demand. In practice, demand-based trading takes many forms. It might take the form of high-frequency trading, where information about an imminent trade is used to trade before the new price is realized. It could be in mining tweets or Facebook posts to gauge sentiment. Extraction could take the form of “partnering,” a practice where brokers sell their demand information (order flow) to hedge funds, who systematically trade against what are presumed to be uninformed traders. Finally, it may mean looking at price trends, often referred to as technical or statistical analysis, in order to discern what information others may be trading on. What all of these practices have in common is that they are not uncovering original information about the future payoff of an asset. Instead, they use public information, in conjunction with private analysis, to profit from what others already know (or don’t know). We capture this general strategy, while abstracting from many of its details, by allowing investors to observe a signal about the non-informational trades of other traders. This demand signal allows our traders to profit in three ways. (i) They can identify and then trade against uninformed order flow, (ii) they can remove noise from the equilibrium price to uncover more of what others know, or (iii) they can exploit the transitory nature of demand shocks to buy before the price rises and sell before it falls. These three strategies have an equivalent representation in the model and collectively cover many of the ways in which modern investment strategies profit from information technology.

A second significant modeling choice is our use of long-lived assets. In this literature, static models have proven very useful in explaining the many forces and trade-offs in a simple and transparent way. However, when the assumption of one-period-lived assets reverses the predictions of the more realistic dynamic model, the static assumption is no longer appropriate. Long-run growth means not only more data processing today, but even more tomorrow. In many instances, the increase today and the further increase tomorrow have competing effects. That competition is a central theme of this paper. Without the long-lived asset assumption, the long-run balanced growth and stagnating liquidity results would be overturned.

Finally, technological progress in its present form gives investors access over time to an increasingly larger set of feasible signals. While there are many possible frameworks that one might use to investigate financial growth, this one makes for a useful lens, because it explains many facts about the evolution of financial analysis, it can forecast future changes that empirical extrapolation alone would miss, and because it offers surprising, logical insights about the financial and real consequences of the structural change. One could go further and argue that some types of data have become relatively easier to collect over time. That may well be true. But changes in relative costs could explain any such pattern. We would not know what results came from relative cost changes and what comes from the fundamental economic forces created by technological change. Our simple problem is designed

to elucidate economic forces, at the expense of many realistic features that could be added.

A. Setup

Investors.—At the start of each date t , a measure-one continuum of overlapping generations investors is born. Each investor i born at date t has constant absolute risk aversion utility over total, end of period t consumption c_{it+1} :

$$(1) \quad U(c_{it+1}) = -e^{-\rho c_{it+1}},$$

where ρ is absolute risk aversion.

Each investor is endowed with an exogenous income that is e_{it} units of consumption goods. At the start of each period, investors decide how much of their income to eat now and how much to spend on a risky asset that pays off at the end of the period.

There is a single tradable, risky asset.³ Its supply is one unit per capita. It is a claim to an infinite stream of dividend payments $\{d_t\}$:

$$(2) \quad d_{t+1} = \mu + G(d_t - \mu) + y_{t+1}.$$

Note, μ and $G < 1$ are known parameters. The innovation $y_{t+1} \sim N(0, \tau_0^{-1})$ is revealed and d_{t+1} is paid out at the end of each period t . While d_{t+1} and d_t both refer to dividends, only d_t is already realized at the start of time t .

In order to disentangle the static and dynamic results, we introduce parameter $\pi \in \{0, 1\}$. When $\pi = 1$, a time- t asset pays $p_{t+1} + d_{t+1}$, the future price of the long-lived asset plus its dividend. When $\pi = 0$, the asset is not long-lived. Its payoff is only the dividend, d_{t+1} . We call the $\pi = 0$ model the “static” model because current information choices do not depend on future or past choices. It is a repeated static problem with an information constraint that changes over time.

In the dynamic ($\pi = 1$) model, an investor born at date t collects dividends d_{t+1} per share and sells the asset at price p_{t+1} to the $t + 1$ generation of investors. In both versions, investors combine the proceeds from risky assets with the endowment that is left $(e_{it} - q_{it}p_t)$ times the rate of time preference $r > 1$, and then consume all those resources. Cohort t consumption can only be realized in $t + 1$, after dividends d_{t+1} are realized and, if $\pi = 1$, the assets are sold to the next cohort. Using c_{it+1} to denote the consumption of cohort t , which takes place in $t + 1$, avoids the double subscript in $c_{t+1,t}$. Thus, the cohort- t investor’s budget constraint is

$$(3) \quad c_{it+1} = r(e_{it} - q_{it}p_t) + q_{it}(\pi p_{t+1} + d_{t+1}),$$

where q_{it} is the shares of the risky asset that investor i purchases at time t and d_{t+1} are the dividends paid out at the end of period t . Since we do not prohibit $c_t < 0$, all pledges to pay income for risky assets are riskless.

³We describe a market with a single risky asset because our main effects do not require multiple assets.

Demand Shocks.—The economy is also populated by a unit measure of noise traders in each period. These traders trade for non-informational reasons.⁴ Each noise trader sells x_{t+1} shares of the asset, where $x_{t+1} \sim N(0, \tau_x^{-1})$ is independent of other shocks in the model. The value of x_{t+1} is revealed at the end of period t . For information to have value, prices must not perfectly aggregate the asset payoff information. This is our source of noise in prices. Equivalently, x_{t+1} could also be interpreted as sentiment. For now, we assume that x_{t+1} is independent over time. In Section IIC, we discuss the possibility of autocorrelated x_{t+1} .

Information Choice.—If we want to examine how the nature of financial analysis has changed over time, we need to have at least two types of analysis between which we can choose. Financial analysis in this model means signal acquisition. Our constraint on acquisition could represent the limited research time for uncovering new information. But it could also represent the time required to process and compute optimal trades based on information that is readily available from public sources.

Investors choose how much information to acquire or process about the end-of-period dividend innovation y_{t+1} , and also about the noisy demand shocks, x_{t+1} . We call $\eta_{fit} = y_{t+1} + \tilde{\epsilon}_{fit}$ a fundamental signal and $\eta_{xit} = x_{t+1} + \tilde{\epsilon}_{xit}$ a demand signal. What investors are choosing is the precision of these signals. In other words, if the signal errors are distributed $\tilde{\epsilon}_{fit} \sim N(0, \Omega_{fit}^{-1})$ and $\tilde{\epsilon}_{xit} \sim N(0, \Omega_{xit}^{-1})$, then the precisions Ω_{fit} and Ω_{xit} are choice variables for investor i .

Next, we recursively define two information sets. The first is all the variables that are known to agent i at the end of period $t - 1$. This information set is $\{\mathbf{I}_{t-1}, y_t, d_t, x_t\} \equiv \mathbf{I}_t^-$. This is what investors of generation t know when they choose what signals to acquire. The second information set is $\{\mathbf{I}_t^-, \eta_{fit}, \eta_{xit}, p_t\} \equiv \mathcal{I}_{it}$. This includes the two signals that the investor chooses to see and the information contained in the equilibrium prices. This is the information set the investor has when they make an investment decision. Let $\mathbf{I}_t = \cup_i \mathcal{I}_{it}$. The time-0 information set includes the entire sequence of information capacity: $\mathbf{I}_0 \supset \{K_t\}_{t=0}^\infty$.

When choosing information ($\Omega_{fit} \geq 0$ and $\Omega_{xit} \geq 0$), investors maximize

$$(4) \quad E[U(c_{it+1}) | \mathbf{I}_t^-],$$

subject to

$$(5) \quad \Omega_{fit}^2 + \chi_x \Omega_{xit}^2 \leq K_t.$$

Here, χ_x is a marginal processing cost parameter that allows us to consider demand data that is easier or harder to process than is fundamental data. Data constraint (5) represents the idea that getting increasingly precise information about a given variable is increasingly difficult.⁵ But acquiring information about a different variable is a separate task, with a shadow cost that is additive.

⁴In previous versions, we microfounded heterogeneous investor hedging demand that would rationalize this trading behavior.

⁵See the online Appendix for results with linear and entropy-based information constraints.

The main force in the model is technological progress in information analysis. Specifically, we assume that K_t is a deterministic, increasing process.

Equilibrium.—An equilibrium is a sequence of investors’ information choices $\{\Omega_{fit}\}$, $\{\Omega_{xit}\}$, prices $\{p_t\}$, and portfolio choices $\{q_{it}\}$, such that:

- (i) Investors choose signal precisions Ω_{fit} and Ω_{xit} to maximize (4), taking the choices of other agents as given. This choice is subject to (5), $\Omega_{fit} \geq 0$, and $\Omega_{xit} \geq 0$.
- (ii) Agents use Bayes’ law to combine prior information \mathbf{I}_t^- with signals η_{fit} , η_{xit} , and p_t , in \mathcal{I}_{it} , and to update beliefs.
- (iii) Investors choose their risky asset investment q_{it} to maximize $E[U(c_{it+1}) | \mathcal{I}_{it}]$, taking the asset price and the actions of other agents as given, subject to the budget constraint (3).
- (iv) At each date t , the risky asset price equates demand, minus demand shocks (sales) and one unit of supply:

$$(6) \quad \int_i q_{it} di - x_{t+1} = 1 \quad \forall t.$$

B. Solving the Model

There are four main steps to solving the model.

Step 1: Solve for the optimal portfolios, given information sets. Each investor i at date t chooses a number of shares q_{it} of the risky asset to maximize the expected utility (1), subject to budget constraint (3). The first-order condition of that problem is

$$(7) \quad q_{it} = \frac{E[\pi p_{t+1} + d_{t+1} | \mathcal{I}_{it}] - r p_t}{\rho \text{var}[\pi p_{t+1} + d_{t+1} | \mathcal{I}_{it}]}.$$

When using the term “investor,” we do not include demand shocks (noise trades).

Step 2: Clear the asset market. Let $\bar{\mathcal{I}}_t$ denote the average investor’s information set, with average signal realizations and average precision. Given the optimal investment choice, we can impose market clearing (6) and obtain a price function that is linear in past dividends d_t , the t -period dividend innovation y_{t+1} , and the demand shock x_{t+1} :

$$(8) \quad p_t = A_t + B(d_t - \mu) + C_t y_{t+1} + D_t x_{t+1},$$

where A_t governs the price level, B is the time-invariant effect of past dividends, C_t governs the information content of prices about current dividend innovations (price information sensitivity), and D_t regulates the amount of demand noise in prices:

$$(9) \quad A_t = \frac{1}{\bar{r}} [\pi A_{t+1} + \mu - \rho \text{var}[\pi p_{t+1} + d_{t+1} | \bar{\mathcal{I}}_t]],$$

$$(10) \quad B = \frac{G}{r - \pi G},$$

$$(11) \quad C_t = \frac{1}{r - \pi G} (1 - \tau_0 \text{var}[y_{t+1} | \bar{\mathcal{I}}_t]),$$

$$(12) \quad rD_t = -\rho \text{var}[\pi p_{t+1} + d_{t+1} | \bar{\mathcal{I}}_t] + \frac{r}{r - \pi G} \text{var}[y_{t+1} | \bar{\mathcal{I}}_t] \frac{C_t}{D_t} \tau_x.$$

The posterior uncertainty about next-period dividend innovations is

$$(13) \quad \text{var}[y_{t+1} | \bar{\mathcal{I}}_t] = (\tau_0 + \bar{\Omega}_{ft} + \bar{\Omega}_{pt})^{-1},$$

where $\bar{\Omega}_{ft} = \int_i \Omega_{fit} di$ is the average fundamental signal precision and $\bar{\Omega}_{pt}$ is the average precision of the information about d_{t+1} , extracted jointly from prices and demand signals. The resulting uncertainty about the future payoff is

$$(14) \quad \text{var}[\pi p_{t+1} + d_{t+1} | \bar{\mathcal{I}}_t] = \pi(C_{t+1}^2 \tau_0^{-1} + D_{t+1}^2 \tau_x^{-1}) + (1 + \pi B)^2 \text{var}[y_{t+1} | \bar{\mathcal{I}}_t].$$

Step 3: Compute the ex ante expected utility. When choosing what information to observe, investors do not know what the signal realizations will be, nor do they know what the equilibrium price will be. The relevant information set for this information choice is \mathbf{I}_t^- .

We substitute optimal portfolio choice (7) and equilibrium price rule (8) into utility (1), and take the beginning of time- t expectation, $-E[\exp(-\rho c_{it+1}) | \mathbf{I}_t^-] = -E[E[\exp(-\rho c_{it+1}) | \eta_{fit}, \eta_{xit}, p_t] | \mathbf{I}_t^-]$, where the equality follows from the law of iterated expectations. Appendix Section A shows that the interim conditional expectation solution is similar to most CARA-normal models: $g_{it} \exp\{(1/2) w^2 \text{var}[\pi p_{t+1} + d_{t+1} | \bar{\mathcal{I}}_t]^{-1}\}$, where g_{it} is a scaling factor, related to i 's endowment, and w is a function of the equilibrium pricing coefficients and model parameters, all of which the investor knows or deduces from the environment.

The key feature of this solution is that the agent's choice variables, Ω_{fit} and Ω_{xit} , show up only through the conditional precision of payoffs, $\text{var}[\pi p_{t+1} + d_{t+1} | \bar{\mathcal{I}}_t]^{-1}$. The reason they only appear there is that the first variance term in asset demand, $\text{var}[\pi p_{t+1} + d_{t+1} | \mathcal{I}_{it}]$, and p_t have ex ante expected values that do not depend on the precision of any given investor's information choices. In other words, choosing to get more data of either type does not, by itself, lead one to believe that payoffs or prices will be particularly high or low. So, information choices amount to minimizing payoff variance $\text{var}[\pi p_{t+1} + d_{t+1} | \mathcal{I}_{it}]$, subject to the data constraint. The payoff variance has terms that the investor takes as given, plus a term that depends on the dividend variance, $\text{var}[y_{t+1} | \mathcal{I}_{it}]$.

So, the information choice problem boils down to this question: what information minimizes dividend uncertainty $\text{var}[y_{t+1} | \mathcal{I}_{it}]$? According to Bayes' law, $\text{var}[y_{t+1} | \mathcal{I}_{it}]$ depends on the sum of fundamental precision Ω_{fit} and price information precision Ω_{pit} (equation (13)). Since price precision is $\Omega_{pit} = (C_t/D_t)^2 (\tau_x + \Omega_{xit})$, the expected utility is a deterministic, increasing function of the sum of Ω_{fit} and $(C_t/D_t)^2 \Omega_{xit}$.

Thus, the optimal information choices maximize the weighted sum of the fundamental and demand precisions:

$$(15) \quad \max_{\Omega_{fit}, \Omega_{xit}} \Omega_{fit} + \left(\frac{C_t}{D_t}\right)^2 \Omega_{xit},$$

subject to

$$\Omega_{fit}^2 + \chi_x \Omega_{xit}^2 \leq K_t, \quad \Omega_{fit} \geq 0, \quad \text{and} \quad \Omega_{xit} \geq 0.$$

The reason we combine fundamental and demand information in this way is because of the linear price equation (8) and Bayes' law. This would be true in any information choice problem, where the objective is a decreasing function of the dividend or payoff uncertainty and prices take the form that they do in (8). Appendix Section A shows that the same information objectives arise with a different utility function, where investors prefer the *early resolution* of uncertainty.

Step 4: Solve for information choices. The first-order conditions yield

$$(16) \quad \Omega_{xit} = \frac{1}{\chi_x} \left(\frac{C_t}{D_t}\right)^2 \Omega_{fit}.$$

This solution implies that information choices are symmetric across agents. Therefore in what follows, we drop the *i* subscript to denote an agent's data processing choice. Moreover, the information set of the average investor is the same as each investor's information set, $\bar{\mathcal{I}}_t = \mathcal{I}_{it}$.

The information choices are a function of the pricing coefficients, like *C* and *D*, which are, in turn, functions of information choices. To determine the evolution of analysis and its effect on asset markets, we need to compute a fixed point to a highly nonlinear set of equations. After substituting in the first-order conditions for Ω_{fit} and Ω_{xit} , we can write the problem as a recursive nonlinear equation in one unknown.

Since this is an overlapping generations model, we expect there to be multiple equilibria. For some parameter values, there are multiple real solutions to this problem. While in some models, multiple equilibria can complicate predictions, in our model, they are not problematic for three reasons: (i) the calibrated model has a unique solution; (ii) the theoretical results hold for any equilibrium; and (iii) in the static model, there is a clear selection criterion. One equilibrium typically converges to the $\Omega_{xit} = 0$ solution, as demand data becomes scarce. The other has price coefficients that become infinite. The obvious choice is the solution with a continuous limit.

C. Interpreting Demand Data Trading

Why are demand signals useful? They do not predict future dividends or future prices. They only provide information about current demand. However, this information is valuable because it tells the investor something about the difference between price and expected asset value. One can see this by looking at the signal extracted from price. Price is a noisy signal about dividends. To extract the price signal, we

subtract the expected value of all the terms besides the dividend and divide by the dividend coefficient C_t . The resulting signal extracted from prices is

$$(17) \quad \frac{(p_t - A_t - B(d_t - \mu) - D_t E[x_{t+1} | \bar{I}_t])}{C_t} = y_{t+1} + \underbrace{\frac{D_t}{C_t} (x_{t+1} - E[x_{t+1} | \bar{I}_t])}_{\text{signal noise}}.$$

Notice how demand shocks x_{t+1} are the noise in the price signal. So information about this demand reduces noises in the price signal. In this way, the demand signal can be used to better extract others' dividend information from the price. In this sense, demand analysis is information extraction.

Of course, real demand traders are not taking their orders then inverting an equilibrium pricing model to infer future dividends. But another way to interpret the demand trading strategy is that it identifies non-informational trades to trade against. In equation (17), when x_{t+1} is high, noise trades are mostly sales and, as $D_t < 0$, prices are low. Moreover, $(D_t/C_t) < 0$ implies that high x_{t+1} makes the expected dividend minus price high, which induces those with demand information to buy. Thus, demand trading amounts to finding the non-informational trades and then systematically taking the opposite side: buying when prices are low and selling when they are high for non-dividend-related reasons. This strategy of trading against uninformed trades is commonly referred to as trading against "dumb money." An alternative way of interpreting the choice between fundamental and demand data is that agents are choosing between decoding private or public signals. Fundamental signals have noise that is independent across agents. These are private. But although demand data's noise is independent, such data are used in conjunction with the price, which is a public signal. The resulting inference about shock y_{t+1} , conditional on the price and the x_{t+1} signal, is conditionally correlated across agents, just as a public signal would be.

The key to the main results that follow is that reducing the noise in x_{t+1} reduces the price noise variance in proportion to $(D_t/C_t)^2$. Put conversely, increasing the precision of information about x_{t+1} (the reciprocal of variance) increases the precision of the dividend information, in proportion to $(C_t/D_t)^2$. The long-run shifts are caused by the marginal rate of substitution of demand signal precision for fundamental signal precision, $(C_t/D_t)^2$, which changes as technology grows.

If we interpret demand trading as finding dumb money, it is easy to see why it would become more valuable over time. If there is very little information, then everyone is "dumb," and finding dumb money is pointless. But when traders are sufficiently informed, distinguishing dumb from smart money, before taking the other side of a trade, becomes essential.

D. Measuring Financial Market Efficiency

To study the effects of financial technology on market efficiency, we assess efficiency in two ways. One efficiency measure is price informativeness. The asset price is informative about the unknown future dividend innovation y_{t+1} . The coefficient C_t on the dividend innovation y_{t+1} in equilibrium price equation (8) governs the extent to which price reacts to a dividend innovation. The coefficient D_t governs

the extent to which demand shocks add noise to the price. Therefore, C_t/D_t is a signal-to-noise ratio that we use to measure price informativeness. It corresponds closely to the price informativeness measure of Bai, Philippon, and Savov (2016).

The other measure of market efficiency is liquidity. Liquidity is the price impact of an uninformed noise trade. That impact is the price coefficient D_t . Note that D_t is negative because it represents selling pressure; the reduced demand lowers the price. So, a more negative D_t represents a higher price impact and a less liquid market. Increasing (less negative) D_t indicates an improvement in liquidity.

E. Existence

One issue with the static ($\pi = 0$) model is that, for any set of parameters, if $K > (\rho^2/2)\sqrt{\chi_x}$, there is no solution to the model. Since one of our main points is to understand what happens as technology grows, this equilibrium nonexistence at high levels of technology is particularly problematic. A key reason to use a model of long-lived assets is that, as long as $\tau_0\tau_x \geq (4\rho/(r(r-G)))^2$, equilibrium exists at every level of K (see equation (104) in the online Appendix). Therefore in what follows, whenever we assume $\pi = 1$, we also assume that τ_0, τ_x are sufficiently large that the existence condition holds. This allows us to explore information choices both when information is scarce and when it is abundant.

The reason equilibrium is preserved is that the unlearnable risk, introduced by future price fluctuations that cannot be known today, keeps prices from being too informative. Because the unlearnable risk grows as technology progresses, the asset never becomes nearly riskless and demand for it never explodes.

For the static results that follow, we ensure existence by assuming that whenever $\pi = 0$, information is not too abundant: $K \leq (\rho^2/2)\sqrt{\chi_x}$.

II. Main Results: A Long-Run Shift in Financial Analysis

This section explores the logical consequences of growth in information (data) processing technology, for financial analysis, trading strategies, and market efficiency. In order to understand what forces produce these results, we first explore the static trade-offs involved in processing either fundamental or demand data. In Section IIA, we consider the effect of an incremental technological change in a setting where an asset's payoff is only its exogenous dividend. When the payoff is the sole exogenous dividend, it clarifies the trade-off between fundamental and demand data and the static forces by which the trade-off is shaped. When $\pi = 0$, future choices or outcomes have no bearing on today's decisions. This is obviously false: by its nature, equity is a long-lived claim. But this setting allows us to clearly derive forces also present in the dynamic model and to distinguish between static and dynamic forces.

The main results center around the model's dynamics. When assets are long-lived ($\pi = 1$, Section IIB), future information risk arises. The risk posed by shocks that will be realized in the future governs the long-run market convergence. We find that as data technology becomes more and more productive, fundamental and demand data processing grow proportionately, price informativeness is high, and there are competing forces in liquidity. For asset prices, the presence of two types

of information matters for the transition path. Demand data also change the value of the price coefficients to which the economy converges in the long run. But, surprisingly, the presence and the growth of demand data does not qualitatively change the long-run price convergence.

A. Short Run Data Trade-Offs

This section investigates the within-period trade-offs in our model. First, we explore what happens in the neighborhood near no information processing, $K \approx 0$. We show that all investors prefer to acquire only fundamental information in this region. Thus, at the start of the growth trajectory, investors primarily investigate firm fundamentals. Next, we prove that an increase in aggregate information processing increases the value of demand information relative to fundamental information. Fundamental information has diminishing relative returns. But in some regions, demand information has increasing returns. What do these phenomena mean for the evolution of analysis? The economy begins by doing fundamental analysis before rapidly shifting to demand analysis. We explore this mechanism, as well as its market efficiency effects, in the following propositions.

To understand why investors with little information capacity use it all on fundamental information, we start by thinking about what makes each type of information valuable. Fundamental information is valuable because it informs an investor about whether the asset is likely to have a high dividend payoff tomorrow. Since prices are linked to current dividends, this also predicts a high asset price tomorrow and thus a high return. Knowing this allows the investor to buy more of the asset when its return will be high and less of it when the return is likely to be low.

In contrast, demand information is not directly relevant to the future payoff or price. But one can still profit from trading on demand. An investor who knows that noisy demands are high will systematically profit by selling (buying) the asset when high (low) demand makes the price higher (lower) than the fundamental value, on average. In other words, demand signals allow one to trade against dumb money.

The next result proves that if the price has very little information embedded in it because information is scarce (K_t is low), then getting demand data to extract price information is not very valuable. In other words, if all trades are “dumb,” then identifying the uninformed trades has no value.

Result 1 (When Information Is Scarce, Demand Analysis Has Zero Marginal Value: Dynamic or Static): As $K_t \rightarrow 0$, for $\pi = 0$ or 1, $dU_1/d\Omega_{xt} \rightarrow 0$.

The proof in online Appendix Section B, which holds for the static and dynamic models ($\pi = 0$ or 1), establishes two key claims: (i) that when $K \approx 0$, there is no information in the price: $C_t = 0$ and (ii) that the marginal rate of substitution of demand information for fundamental information is proportional to $(C_t/D_t)^2$. In particular, $dU_1/d\Omega_{xt} = (C/D)^2 dU_1/d\Omega_{ft}$. Thus, when the price contains no information about future dividends ($C_t = 0$), then analyzing demand has no marginal value ($(C_t/D_t)^2 = 0$). Demand data are only valuable in conjunction with the current price, p_t , because it allows one to extract more information from the price. Demand data trading when $K_t = 0$ is like removing noise from a signal that has

no information content. Put differently, when there is no fundamental information, the price perfectly reveals noise trading. There is no need to process data on noisy demand if it can be perfectly inferred from the price.

This result explains why analysts focus on fundamentals when financial analysis productivity is low. In contrast, when prices are highly informative, demand information is like gold because it allows one to exactly identify the price fluctuations that are not informative and are therefore profitable to trade on. The next results explain why demand analysis increases with productivity growth and why it may eventually start to crowd out fundamental analysis.

Next we turn to understanding how technological growth affects prices. Specifically, we examine the static effect on price informativeness, C_t/D_t (the signal-to-noise-ratio); the price information sensitivity, C_t ; and (il)liquidity, D_t . Technology improvements affect the price coefficients through two channels: a change in fundamental analysis and a change in demand analysis. Using the chain rule, we can describe which portion of the total effect of the change in information technology, K_t , works through which channel.

Result 2 (Price Response to Technological Growth: Static): For $\pi = 0$,

- (a) $\frac{dC_t/|D_t|}{dK_t} > 0$;
 $\frac{\Omega_{ft}}{2K_t(2\Omega_{xt}C_t/D_t + \rho)} > 0$ of this effect comes through changes in fundamental information Ω_{ft} ;
 $-\frac{\rho C_t/D_t + \Omega_{ft}}{2K_t(2\Omega_{xt}C_t/D_t + \rho)} > 0$ of this effect comes through changes in demand information Ω_{xt} ;
- (b) $\frac{dC_t}{dK_t} > 0$;
- (c) $\frac{d|D_t|}{dK_t} > 0$ if and only if $K_t < \bar{K}_D$, where \bar{K}_D is defined by equation (82) in the online Appendix.

For each type of analysis, there is a direct and an indirect effect of K_t . The direct effect is what would arise if we kept the signal-to-noise ratio, i.e., the marginal rate of transformation across the two types of analysis, constant. An additional indirect effect arises because the change in the type of analysis affects the signal-to-noise ratio, which then affects the information choices. However, using the envelope theorem, the marginal indirect effect is zero at the optimum. Therefore, the results above reflect only the direct effect. The proof of Result 5 shows that our finding that more information today increases the concurrent price informativeness also carries over to the dynamic model ($\pi = 1$).

The concern about the deleterious effects of financial technology on market efficiency stems from the worry that technology will deter the research and discovery of new fundamental information. This concern is not unwarranted. Not only does more fundamental information encourage the extraction of information from demand, but once it starts, demand analysis feeds on itself. The next corollary shows

that when $\pi = 0$, aggregate demand analysis increases an individual's incentive to learn about demand.

For most of our results, we use Ω_{xt} to indicate the demand of every investor, because all are symmetric. At this point, it is useful to distinguish between one particular investor's choice, Ω_{xit} , and the aggregate symmetric choice, Ω_{xt} .

Result 3 (Complementarity in Demand Analysis: Static): For $\pi = 0$, $\frac{\partial \Omega_{xit}}{\partial \Omega_{xt}} \geq 0$.

Fundamental information, Ω_{ft} , exhibits strategic substitutability in information, just as in Grossman and Stiglitz (1980). But for demand information, the effect is the opposite. More precise average demand information (higher Ω_{xt}) can increase $(C_t/D_t)^2$, which is the marginal rate of the substitution of demand information for fundamental information. The rise in the relative value of demand data is what makes investors shift their data analysis from fundamental to demand when others do more demand analysis. That is complementarity. It holds in both the static model and the dynamic model, with conditions. (For a result with $\pi = 1$, see online Appendix Section B).

Intuitively, higher signal-to-noise (more informative) prices encourage demand trading because the value of demand analysis comes from the ability to better extract the signal from prices. In this model (as in most information processing problems), it is easier to clear up relatively clear signals than it is very noisy ones. So the aggregate level of demand analysis improves the signal clarity of prices, which makes demand analysis more valuable.

Why does the price signal-to-noise ratio rise? From (11), we know that C_t is proportional to $1 - \tau_0 \text{var}[y_{t+1} | \bar{I}_t]$. As either type of information precision (Ω_{ft} or Ω_{xt}) improves, the uncertainty about the next period's dividend innovation $\text{var}[y_{t+1} | \bar{I}_t]$ declines and C_t increases. The variable D_t is the coefficient on noise x_{t+1} . The price impact of uninformative trades $|D_t|$ may also increase with information (Result 2c). But $|D_t|$ does not rise at a rate faster than C_t . Thus, the ratio $C_t/|D_t|$, which is the signal-to-noise ratio of prices, and the marginal value of demand precision, increases with more information.

The final result of this section characterizes how different types of analysis change as there is technological progress, in a world with one-period-lived assets.

Result 4 (Fundamental and Demand Analyses' Response to Technological Growth: Static): For $\pi = 0$,

- (a) fundamental analysis initially grows and then declines, $d\Omega_{ft}/dK_t > 0$ if and only if $K_t < \bar{K}_f = (\sqrt{3}/2)\bar{K}$,
- (b) demand analysis is monotonically increasing, $d\Omega_{xt}/dK_t > 0$.

As technology improves, both types of information analysis initially grow. However, fundamental analysis experiences two competing forces. On one hand, more available capacity increases fundamental analysis. On the other hand, the higher marginal rate of substitution between demand and fundamental analyses (a higher signal-to-noise ratio) dampens the level of fundamental analysis. When there is

not much information available, the first force dominates. Once the information processing capacity grows beyond a certain threshold, substitution toward demand analysis takes over and fundamental analysis falls. The online Appendix shows how this intuition also carries over to the dynamic economy ($\pi = 1$). We characterize the conditions under which price informativeness, price information sensitivity, and liquidity improve in response to exogenous changes in data endowments (Result 8).

A rise in data processing efficiency is not the only force that can boost price informativeness. Online Appendix Section C also derives comparative statics. The results show that greater risk tolerance (lower ρ) or an increase in demand data processing efficiency (lower χ_x) both raise $C_t/|D_t|$. Trends in either variable could therefore also prompt a switch to demand-based trading strategies.

B. Dynamic Results

Next, we explore the model's dynamic forces. By assuming $\pi = 1$, the asset's payoff is $p_{t+1} + d_{t+1}$, as with a traditional equity payoff. This introduces one new concept, future information risk: knowing that tomorrow's investors will get more information makes it harder to predict what those investors will believe, which makes future demand and future prices more uncertain. Since the future price is part of the payoff to today's asset, future information risk makes this payoff more uncertain. Assets look riskier to investors. The result that future information creates risk is central to the main finding of long-run balanced data processing. Without long-lived assets, information learned tomorrow cannot affect the payoff risk today. Long-lived assets are integral to all the results that depend on future information risk, including our main result, the long-run balanced growth of data processing.

Future information risk is the part of payoff risk $\text{var}[p_{t+1} + d_{t+1} | \bar{\mathcal{I}}_t]$ that comes from $\{K_{t+1}, \dots, K_\infty\} > 0$. This risk arises from shocks to tomorrow's price that are unknowable today. But not all unknowable shocks come from future information. Tomorrow's demand shock will create price noise, regardless of whether others learn about it. Instead, future information grows the impact of unknowable price shocks. The next result shows that more learning tomorrow makes tomorrow's prices, p_{t+1} , more sensitive to shocks that are unlearnable today, and this additional risk makes today's prices less informative.

Result 5 (Future Information Creates Risk, Reduces Informativeness: Dynamic Only): If $\pi = 1$ and ρ is sufficiently low, then an increase in future information K_{t+1} ,

- (a) increases payoff risk $\text{var}[p_{t+1} + d_{t+1} | \bar{\mathcal{I}}_t]$, and
- (b) reduces current price informativeness $C_t/|D_t|$.

The proof in online Appendix Section B goes further and makes it clear that what reduces informativeness is the fact that the information will be learned in the future. More information today (higher K_t) unambiguously increases price informativeness $C_t/|D_t|$ today.

Conceptually, future information creates risk because, if many investors will trade on precise $(t + 1)$ information tomorrow, then tomorrow's price will be very sensitive to the next day's dividend information y_{t+2} and possibly to the next day's demand information x_{t+2} . But investors today do not know what will be learned tomorrow. Therefore, tomorrow's analysis makes tomorrow's price (p_{t+1}) more sensitive to shocks about which today's investors are uninformed. Because future information has no effect on today's dividend uncertainty, $\text{var}[y_{t+1} | \bar{\mathcal{I}}_t]$, and because it raises future price uncertainty, the net effect of future information is to raise today's payoff variance. That is what creates risk today.

Mathematically, the relationship between tomorrow's price coefficients and future information risk is evident in the C_{t+1} and D_{t+1} coefficients in the formula for future information risk. Future information affects today's risk because the payoff variance is $\text{var}[p_{t+1} + d_{t+1} | \bar{\mathcal{I}}_t] = C_{t+1}^2 \tau_0^{-1} + D_{t+1}^2 \tau_x^{-1} + (1 + B)^2 \text{var}[y_{t+1} | \bar{\mathcal{I}}_t]$. The key terms are $C_{t+1}^2 \tau_0^{-1} + D_{t+1}^2 \tau_x^{-1}$. Future information affects the future price coefficients C_{t+1}^2 and D_{t+1}^2 . We know that time- t information increases period- t information content C_t . Similarly, time $t + 1$ information increases C_{t+1} . Future information may increase or decrease D_{t+1} . But as long as $C_{t+1}/|D_{t+1}|$ is large enough, the net effect of $t + 1$ information is to increase $C_{t+1}^2 \tau_0^{-1} + D_{t+1}^2 \tau_x^{-1}$.

One reason future information risk is important is that it can reduce today's liquidity. It makes future price p_{t+1} more sensitive to future information and thus harder to forecast today. That raises the time- t asset payoff risk $\text{var}[p_{t+1} + d_{t+1} | \bar{\mathcal{I}}_t]$. A riskier asset has a less liquid market. We can see this relationship in the formula for D_t (equation (12)), where $\text{var}[p_{t+1} + d_{t+1} | \bar{\mathcal{I}}_t]$ shows up in the first term. Thus, future information reduces today's liquidity.

Technology growth improves information today and then improves it again tomorrow. That means the static and dynamic effects are in competition. The net effect of the two is sometimes positive, sometimes negative. Price volatility does not become arbitrarily large. Just as the conditional variance of dividends converges to zero, price volatility converges to a positive, finite level. The point is that the net effect is not as clear-cut as what a static information model predicts. We learn that with long-lived assets, information technology efficiency and liquidity are not synonymous. In fact, because financial technology makes future prices more informative, it can also make markets function in a less liquid way.

The static result that demand analysis feeds on itself suggests that in the long run, it will completely crowd out fundamental analysis. But that does not happen. When demand precision (Ω_{xt}) is high, the conditions for Proposition 3 break down. The next result tells us that in the long run as information becomes abundant, growth in fundamental and demand analyses becomes balanced. This result for the long-lived asset contrasts with the static asset Result 4, where fundamental and demand analyses diverge.

Result 6 (High-Information Limit: Dynamic Only): If $\pi = 1$ and $K_t \rightarrow \infty$, both analysis choices Ω_{ft} and Ω_{xt} tend to ∞ such that

- (a) Ω_{ft}/Ω_{xt} does not converge to 0;
- (b) Ω_{ft}/Ω_{xt} does not converge to ∞ ; and

- (c) if τ_0 is sufficiently large, there exists an equilibrium where Ω_{ft}/Ω_{xt} converges to a finite, positive constant.
- (d) No perfect liquidity: There is no equilibrium for any date t with $D_t = 0$.

See online Appendix Section B for the proof and an expression (104) for the lower bound on τ_0 .

It is not surprising that fundamental analysis will not push demand analysis to zero (part (a)). We know that more fundamental analysis lowers the value of additional fundamental analysis and raises the value of demand analysis by increasing $C_t/|D_t|$. This is the force that prompts demand analysis to explode at lower levels of information K .

But what force restrains the growth of demand analysis? It's the same force that keeps liquidity in check: information today, competing with the risk of future information that will be learned tomorrow. The first-order condition tells us that the ratio of fundamental to demand analyses is proportional to the squared signal-to-noise ratio, $(C_t/D_t)^2$. If this ratio converges to a constant, the two types of analysis remain in fixed proportion. Recall from Result 5 that information acquired tomorrow reduces D_t . That is, D_t becomes more negative, but larger in absolute value. On the other hand, as data observed today become more abundant, the price information sensitivity (C_t) grows and liquidity improves: D_t falls in absolute value. As data processing grows, the upward force of current information and the downward force of future information bring C_t and D_t each to rest at a constant, finite limit. For C_t , this infinite-data limit is $\bar{C} = 1/(r - G)$. For D_t , the limit is $\bar{D} = \left[-r \pm \sqrt{r^2 - 4(\rho/(r - G))^2 \tau_0^{-1} \tau_x^{-1}} \right] / (2\rho\tau_x^{-1})$. Lemma 4 in the online Appendix explores the properties of the ratio $C_t/|D_t|$ in this limit. It shows formally that $C_t/|D_t|$ is bounded above by the inverse of future information risk. When assets are not long-lived, their payoffs are exogenous, future information risk is zero, and $C_t/|D_t|$ can grow without bound. Put differently, without a long-lived asset, the limit on $C_t/|D_t|$ is infinite. Data processing would not be balanced.⁶

C. Persistent Demand or Information about Future Events

A key to many of our results is that the growth of financial technology creates more and more future information risk. This risk arises because shocks that affect tomorrow's prices are not learnable today, which raises the question: what if information about future dividend or demand shocks was available today?

As long as there is still some uncertainty and thus something to be learned in the future, future information will still create risk for returns today. Tomorrow's price would depend on information learned tomorrow about shocks that will materialize in $t + 2$ or $t + 3$. That new information, observed in $t + 1$, will affect $t + 1$ prices.

⁶Another way to introduce unlearnable risk is to simply have a static model with some portion of the payoff that cannot be learned. While this economy would have an equilibrium and analysis would converge, the convergence path and the ratio of analysis the economy converges to would differ. A version of the long-lived asset effect can also arise in a dynamic model with only fundamental analysis (see Cai 2016).

The general point is this: new information is constantly arriving; it creates risk. Whether it is about tomorrow, the next day, or the far future, this yet-to-arrive information will have an uncertain impact on future prices. When information processing technology is poor, the poorly processed future information has little price impact. Thus, scarce future information poses little risk. When information processing improves, the risk of unknown future information grows.

Of course, if persistent demand was the source of future information, then signals about demand, x_{t+1} , would be payoff relevant. The x_{t+1} signal would be informative about x_{t+2} , which affects the price p_{t+1} and thus the payoff of a time t risky asset. In such a model, agents would attempt to distinguish the persistent and transitory components of demand. The persistent, payoff-relevant component would play the role of dividend information in this model. The transitory component of demand would play the role of the i.i.d. x_{t+1} shock.

A broader interpretation of the model is that there are three categories of activities: learning about persistent payoff-relevant shocks, learning about transitory price shocks, and future learning about any shocks that cannot be known today. Exactly what is in each of these categories is a matter of interpretation.

D. Price Informativeness, Liquidity, Welfare, and Real Economic Output

Why are price informativeness and liquidity sensible measures of market efficiency? In this setting, all dividends are exogenous. Information has no effect on firms' real output. It only facilitates reallocation of these goods from one trader to another. In fact, the social optimum is achieved with full risk-sharing, which arises when there are no data and beliefs are therefore symmetric.

Does all this mean that data are bad for society? Not necessarily. One way to approach social welfare is to consider maximizing price informativeness to be a social objective. The question then becomes: are data choices socially efficient? It turns out that they do indeed maximize price informativeness.

Result 7 (Social Efficiency: Static): For $\pi = 0$, the equilibrium attention allocation maximizes price informativeness $C_t/|D_t|$.

The reason that social and private objectives are aligned is that the same sufficient statistic governs both. Individual data choices that maximize expected utility do so by maximizing the precision of information about dividend innovations: $\text{var}[y_{t+1}|\mathcal{I}_t]^{-1}$. For the representative investor, price informativeness also depends on aggregate data choices, but only through $\text{var}[y_{t+1}|\mathcal{I}_t]^{-1}$. Price informativeness is increasing in $\text{var}[y_{t+1}|\mathcal{I}_t]^{-1}$. Thus the data choices that maximize individual utility maximize $\text{var}[y_{t+1}|\mathcal{I}_t]^{-1}$, and therefore also price informativeness.

An alternative approach is to relate financial markets and the real economy. This model is stripped to its barest essentials to make its logic most transparent. For that reason, it has no link between financial and real economic outcomes. With that link added back in, liquid and information-rich financial markets can have real benefits.

Online Appendix Sections C.5 and C.6 sketch two models of real economic production where the amount produced by a firm depends on price information sensitivity or liquidity, as in David, Hopenhayn, and Venkateswaran (2016). In the

first model, a manager makes a costly effort to increase the future value of the firm and is compensated with equity. When the equity price is more informative, that means the price reacts more to effort and the associated output. That makes equity a better incentive for providing effort and raises firm output. In the second model, a firm needs to issue equity to raise funds for additional real investment. When markets are illiquid, issuing equity exerts strong downward pressure on the equity price. This reduces the firm's ability to raise revenue, reduces the size of the capital investment, and depresses output. While these models are just caricatures of well-known effects, they illustrate why the objects that are the central focus of analysis in this paper, price informativeness and liquidity, are of such interest.

III. Illustrating Financial Technology Growth: Numerical Example

Our main results reveal competing forces are work. Quantifying the model provides some understanding of which effect is likely to dominate. The equilibrium effects that we focus on are price informativeness and liquidity. A common concern is that as financial technology improves, the extraction of information from demand will crowd out original research, and in so doing, will reduce the informativeness of market prices. On the flip side, if technology allows investors to identify uninformed trades and take the other side of those trades, such activity is thought to improve market liquidity. While each argument has a grain of truth to it, countervailing equilibrium effects mean that neither conjecture is correct.

We begin by quantifying the forces that make demand information more valuable over time. Then we explore why the shift from information production to extraction does not harm price informativeness or improve liquidity. Finally, we ask whether the model contradicts the long-run trends in price volatility, decomposes demand and fundamental data effects, explores alternative parameter values, and considers the possibility of biased technological change.

A. Calibration

Our calibration strategy is to measure the growth of computer processor speed directly to discipline technology K and then to estimate our equilibrium price equation on the recent asset price and dividend data, assuming assets are long-lived. Then we use the time path of the price coefficients and the signal-to-noise ratio to calibrate the model parameters.

We describe the data and then the moments of the data and model that we match to identify the model parameters. Most of these moments come from estimating a version of our price equation (8) and choosing parameters that match the price coefficients in the model with the data. In the next section, we report the results.

Measuring Data Growth.—Investors can acquire information about asset payoffs y_{t+1} or demand x_{t+1} by processing digital data, which is coded in binary code. To calibrate the growth rate of data technology, we choose a growth rate for K , which implies a growth rate of bits processed equal to the growth rate of computer processor speed or cloud computing capacity.

To map the economic measure of data K into a binary string length, we use a concept from information theory called the Gaussian channel. In a Gaussian channel, all data processing is subject to noise (error).⁷ The number of bits required to transmit a message is related to the channel's signal-to-noise ratio. Clearer signals can be transmitted through the channel, but they require more bits. The relationship between bits and signal precision for a Gaussian channel is $bits = 1/2\ln(1 + \text{signal-to-noise})$ (Cover and Thomas 1991, Theorem 10.1.1). The signal-to-noise is the ratio of signal precision to prior precision. In the notation of this model, if the prior precision is τ , the number of bits \tilde{b} required to transmit Ω units of precision in a signal is $\tilde{b} = (1/2)\ln(1 + (\Omega/\tau))$.

If this is true both for fundamental precision Ω_{ff} and for demand precision Ω_{xt} , and presumably, each piece of data is transmitted separately, then the total number of bits processed b is the sum of fundamental and demand bits:

$$(18) \quad b = \frac{1}{2}\ln\left(1 + \frac{\Omega_{ff}}{\tau_0}\right) + \frac{1}{2}\ln\left(1 + \frac{\Omega_{xt}}{\tau_x}\right).$$

Using this mapping, we choose a growth rate of K_t , such that the equilibrium choices of Ω_{ff} and Ω_{xt} imply a growth rate of bits that matches the data. We explain the procedure we use to do so shortly. That path of K_t is

$$(19) \quad K_t = 0.00095 \times 2^{0.49(t-1)} \quad \text{for } t = 1, \dots, T.$$

The multiplier 0.00095 is simply a normalization to keep units from becoming too large. The choice of 0.49 in the exponent ensures that the (cumulative) average bit growth is close to 20 percent per year for the first 18 periods, because these periods correspond to the past observed data in our calibration, as explained later in the section. The 20 percent annual growth rate of bit processing reflects evidence from multiple sources. One source is hardware improvement: the speed of frontier processors has grown by 27 percent since the 1980s, and has more recently slowed to 20 percent growth per year (Hennessy and Patterson 2009). Another fact that supports this rate of growth is the 19 percent growth rate of the workloads in data centers (22 percent for cloud data centers), where most banks process their data (Cisco 2018).

Asset Data.—The data are obtained from Compustat and cover S&P500 firms over the 1985–2015 period.⁸ We construct a measure $ea_{i,t}$ for relevant firm i as earnings before interest and taxes (EBIT) divided by total assets. For each firm i , we fit an AR(1) model to the time series $ea_{i,t}$ and use the residual $ue_{i,t}$ as a measure of $y_{i,t}$. Finally, we construct a measure $lma_{i,t}$, which is the log of market capitalization over total assets.⁹

Our measure of dividends is the lagged value of $ea_{i,t}$ and therefore μ is estimated as the sample mean of $ea_{i,t-1}$ averaged over firms and dates. Then, to create the time

⁷ As Cover and Thomas (1991, p. 240) explains: “The additive noise in such channels may be due to a variety of causes. However, by the central limit theorem, the cumulative effect of a large number of small random effects will be approximately normal, so the Gaussian assumption is valid in a large number of situations.”

⁸ The Compustat database is from Standard and Poor's (1985–2015) and can be accessed via Wharton Research Data Services (WRDS).

⁹ That is, if M is market capitalization and A total assets, $\log(M/A)$.

series of price coefficients, we run the following cross-sectional regression for each year $t = 1, \dots, T$,¹⁰

$$(20) \quad lma_{i,t} = \hat{A}_t + \hat{B}_t(ea_{i,t-1} - \mu) + \hat{C}_t ue_{i,t} + \hat{H}_{i,t} Y_{i,t} + \hat{\epsilon}_{i,t},$$

where $Y_{i,t}$ is a collection of dummies at the SIC3 industry level.¹¹

This first step results in a time series of estimated coefficients that are imperfect measures of A_t, B_t , and C_t in the model. In addition, the squared residuals $\hat{\epsilon}_{i,t}$ correspond to the model's $D_t^2 \tau_x^{-1}$. For \hat{B}_t, \hat{C}_t , and $1/N \sum_{i=1}^N \hat{\epsilon}_{i,t}^2$, we remove their high-frequency time-series fluctuations. To do this, we simply regress each one on a constant, a time trend, and a squared time trend. That is, if the estimated coefficient is x_t , we estimate β_a, β_b , and β_c in $x_t = \beta_a + \beta_b t + \beta_c t^2 + \nu_t$. Then, we construct a smoothed series as $x_{t+1} = \hat{\beta}_a + \hat{\beta}_b t + \hat{\beta}_c t^2$, where $\hat{\beta}_a, \hat{\beta}_b$, and $\hat{\beta}_c$ are the estimates of β_a, β_b , and β_c . This procedure results in three series: \tilde{B}, \tilde{C} , and $\tilde{D}_t^2 \tau_x^{-1}$, which we use to calibrate to the average rate of the coefficient change in the last 30 (approximately) years.¹²

Estimation.—We pick three parameters of the model directly. As described above, μ is set to match the average earnings. The riskless rate $r = 1.05$ is set to match a 5 percent annual net return. Risk aversion clearly matters for the level of the risky asset price, but it is not well identified. Doubling the variance and halving the risk aversion mostly just redefines the units of risk. In what follows, we set $\rho = 0.05$. In the online Appendix, we explore other values to show that our results are robust to changes of this kind.

Four parameters, $\theta = (G, \tau_0, \tau_x, \chi_x)$, as well as the growth rate of K_t , remain to be estimated. The sequence of information capacities, K_t , is chosen to match the average bit growth rate, which depends on the equilibrium outcomes. We use the following procedure: pick a growth rate for K_t and estimate θ as explained below. Given the θ estimate, compute the equilibrium and find the date that corresponds most closely to the most recent observations in the constructed calibration series (2015), \hat{t} .¹³ Next calculate the corresponding average bit growth rate for period $(0, \hat{t})$. Iterate on this procedure until the 20 percent average target bit growth rate is hit.

Finally we estimate $\theta = (G, \tau_0, \tau_x, \chi_x)$. We describe four moments derived from the model. Using these four moment conditions, we estimate the parameter vector, θ , by the generalized method of moments. The theory constrains $\theta \in [0, 1) \times (0, \infty)^3$. Let $X \in \mathbb{R}^{(T-1) \times 5}$ be the sequences $\tilde{B}_t, \tilde{C}_t, \tilde{D}_t^2 \tau_x^{-1}$, and K_t .¹⁴ The first three moments below are the equilibrium solution for the price coefficients B_t, C_t , and D_t . They are simply rearrangements of (44), (45), and (46). The fourth equation uses

¹⁰ See online Appendix for a discussion of the generated regressor.

¹¹ Notice that B_t makes a return here with a t subscript, even though we proved that B was a constant. That is because when we estimate price coefficients in each year, the B_t estimate is, of course, not constant. We could first average B_t over time and use the average. But using the B_t that is consistent with the C_t and D_t for that period t results in more precise parameter estimates.

¹² Allowing just a constant and time trend yields very similar results.

¹³ Specifically, $t = 18$ minimizes $(A_t - a)^2 + (C_t - c)^2 + (D_t - d)^2$ over the 150 simulated model periods, where $a = \tilde{A}_T, c = \tilde{C}_T$, and $d = -\sqrt{(D_T^2 \tau_x^{-1} \times \tau_x)}$.

¹⁴ The last observation has to be dropped, due to the presence of $t + 1$ parameters in the time- t moment conditions.

the information budget constraint and the pricing solutions to characterize the signal-to-noise ratio in prices C_t/D_t . It is a variation of equation (69). Now,

$$(21) \quad g_{1,t}(\theta, r, \rho, X) := \frac{1}{r}(1 + B_{t+1})G - B_t,$$

$$(22) \quad g_{2,t}(\theta, r, \rho, X) := \frac{1}{r - G} \left(1 - \tau_0 \left(\tau_0 + \Omega_{ft} + \left(\frac{C_t}{D_t} \right)^2 \left(\tau_x + \Omega_{ft}/\chi_x \left(\frac{C_t}{D_t} \right)^2 \right) \right)^{-1} \right) - C_t,$$

$$(23) \quad g_{3,t}(\theta, r, \rho, X) := \left(\frac{\tau_x C_t}{r - G D_t} - \frac{\rho r}{(r - G)^2} \right) \times \left(\tau_0 + \Omega_{ft} + \left(\frac{C_t}{D_t} \right)^2 \left(\tau_x + \Omega_{ft}/\chi_x \left(\frac{C_t}{D_t} \right)^2 \right) \right)^{-1} - \frac{\rho}{r} (C_{t+1}^2 \tau_0^{-1} + D_{t+1}^2 \tau_x^{-1}) - D_t,$$

$$(24) \quad g_{4,t}(\theta, r, \rho, X) := \left(\frac{C_t}{D_t} \right)^3 Z_t \tau_x + \frac{C_t}{D_t} \left(\frac{\rho r}{r - G} + Z_t \tau_0 \right) + \left(1 + \frac{C_t}{D_t} Z_t \right) \frac{K_t}{\Omega_{ft}}.$$

Note that Z_t captures future information risk and the equilibrium demand for fundamental information Ω_{ft} comes from combining information capacity constraint (5) with first-order condition (16):

$$(25) \quad Z_t := \frac{\pi \rho}{r} (r - \pi G) (C_{t+1}^2 \tau_0^{-1} + D_{t+1}^2 \tau_x^{-1}),$$

$$(26) \quad \Omega_{ft} := \left(\frac{K_t \chi_x}{\chi_x + (C_t/D_t)^4} \right)^{1/2}.$$

According to the model, each of these four moments $g_{1,t}$ though $g_{4,t}$ should be zero at each date t . To estimate the four parameters $\theta = (G, \tau_0, \tau_x, \chi_x)$ from these four moment equations, we compute the actual value of $g_{1,t}$ though $g_{4,t}$ at each date t , for a candidate set of parameters θ . We average those values (errors) over time, square them, and sum over the four moments. Formally, let $\bar{g}(\theta)$ (a 4×1 vector) be the time-series mean of the $(T - 1) \times 4$ matrix

$$(27) \quad g(\theta) := \begin{bmatrix} g_{1,1}(\theta, r, \rho, X) & g_{2,1}(\theta, r, \rho, X) & g_{3,1}(\theta, r, \rho, X) & g_{4,1}(\theta, r, \rho, X) \\ \vdots & \vdots & \vdots & \vdots \\ g_{1,T-1}(\theta, r, \rho, X) & g_{2,T-1}(\theta, r, \rho, X) & g_{3,T-1}(\theta, r, \rho, X) & g_{4,T-1}(\theta, r, \rho, X) \end{bmatrix},$$

and let I be the 4×4 identity matrix. The estimated parameter vector $\hat{\theta}$ solves

$$\hat{\theta} := \underset{\theta \in (0,1) \times (0,\infty)^3}{\operatorname{argmin}} \bar{g}(\theta)' I \bar{g}(\theta).$$

We optimize with constraints that τ_0 , τ_x , and χ_x are positive and then check that the estimated value of G lies within its admissible range $[0, 1)$. This procedure produces the following parameter values.

TABLE 1—PARAMETERS

ρ	r	μ	G	τ_0	τ_x	χ_x
0.05	1.02	0.04	0.98	80.08	19.75	21.12

Computation and Equilibrium Selection.—We use the estimated parameters to solve for the long-run stable pricing coefficients. Specifically, we set C_T/D_T to the smaller solution to quadratic equation (103) in the online Appendix. Moreover, we choose $T = 150$ to be our long run, and create path K_t according to equation (19). Then we solve the model backward: knowing the time- t price coefficients allows us to numerically find the root of equation (67) to obtain C_{t-1}/D_{t-1} , from which we can solve for the information choices and the price function coefficients in period $t - 1$.

The nonlinear equation in C_t/D_t that characterizes the solution can have multiple solutions. As it happens, for the parameter values that we explore, given a C_{t+1} and D_{t+1} , this equation has only one real root. Moreover, the date that corresponds most closely to the most recent observations in our constructed calibration series, corresponding to data from 2015, is period 18.

B. Result: The Transition from Fundamental to Demand Analysis

With our calibrated model, we can now illustrate the three phases of information processing.

Figure 1 shows that initially demand analysis is scarce. Consistent with Result 1, we see that when the ability to process information is limited, almost all of that ability is allocated to processing fundamental information. Once fundamental information is sufficiently abundant, demand analysis takes off. Not only does demand processing surge, but it increases by so much that the amount of fundamental information declines, even though the total ability to process information has improved. Once demand trading takes off, it quickly comes to dominate fundamentals-based trading. In the long run, the two types of information processing grow at the same rate (albeit at different levels).

C. Price Informativeness

Prices are informative to the extent that they reflect a change in future dividends relative to a change in aggregate supply. Our equilibrium price solution (8) reveals that the absolute value of this relative price impact, $|(dp_t/dy_{t+1})/(dp_t/dx_{t+1})|$ is $C_t/|D_t|$. As the productivity of a financial analysis rises, and more information is acquired and processed, the information sensitivity of price (C_t) rises, while the illiquidity ($|D_t|$) changes nonmonotonically.

Our analysis shows that overall as technology improves, investors are able to better discern an asset's true value, and price informativeness rises. The thick solid line labeled $(C_t/D_t)^2$ in Figure 2 confirms that as financial analysis becomes more productive, informativeness rises. This change is driven primarily by the increase in the information sensitivity of prices C_t . The effect of a one-unit change in the dividend innovation, which is about 2 standard deviations, increases the price by between 0

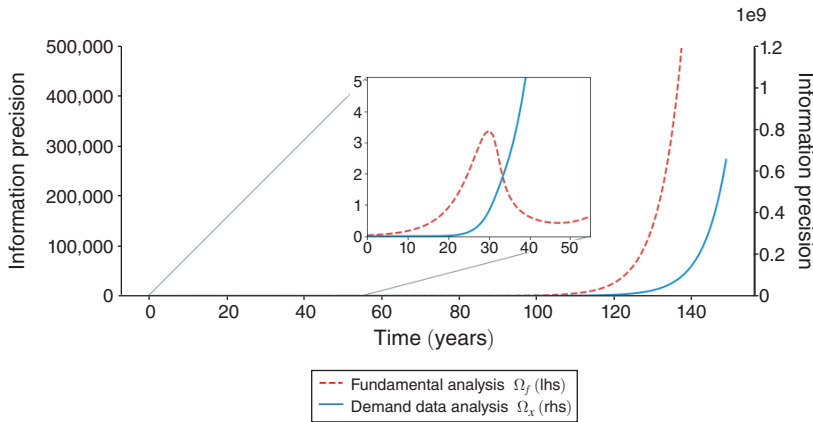


FIGURE 1. EVOLUTION OF FUNDAMENTAL ANALYSIS AND DEMAND ANALYSIS

Notes: What drives the change over time is an increase in total information processing K . Fundamental information is choice variable Ω_f , scaled by fundamental variance τ_0^{-1} . Demand information is choice variable Ω_x , scaled by nonfundamental demand variance τ_x^{-1} .

and 8 units. Because the average price level is about 80, this 2 standard deviation shock to dividends produces a negligible price change for very low levels of technology and a 10 percent price rise when financial technology becomes more advanced.

D. Price Impact of Trades (Illiquidity)

A common metric of market liquidity is the sensitivity of an asset's price to a buy or sell order. In our model, price impact is the impact of a one-unit noise trade $(dp_t/d(-x_{t+1}))$.¹⁵ Linear price solution (8) reveals that price impact is $dp_t/d(-x_{t+1}) = |D_t|$.

Looking at the thin line in Figure 2, we see that the price impact of noise trades, $|D_t|$, rises in the early periods when only Ω_{ft} is increasing. It then declines, but it stays positive as information becomes more abundant. Since 2015 corresponds to period 18, this suggests that we are close to the maximum illiquidity. But surprisingly, the changes are quite small. A noise trade that is the size of 1 percent of all outstanding asset shares would increase the price by 0.05–0.06 units. Since the average price is 80, this amounts to a 0.6–0.7 percent (60–70 basis point) increase in the price. By exploring different parameters, we see that the dynamics of market liquidity can vary. But what is consistent is that the liquidity changes are small, compared to the price information changes.

Flat liquidity is the result of two competing forces. Recall from Section II that the liquidity of a risky asset is determined largely by the riskiness (uncertainty) of its payoff. The static force $(r/(r-G))\text{var}[y_{t+1}|\bar{I}_t](C_t/D_t)$, which reduces risk and

¹⁵We consider a noise trade because the impact of an information-based trade would reflect the fundamental (future dividend), which must have moved in order to change the information. We have already explored the question of how much a change in the fundamental changes the price. That is price information sensitivity.

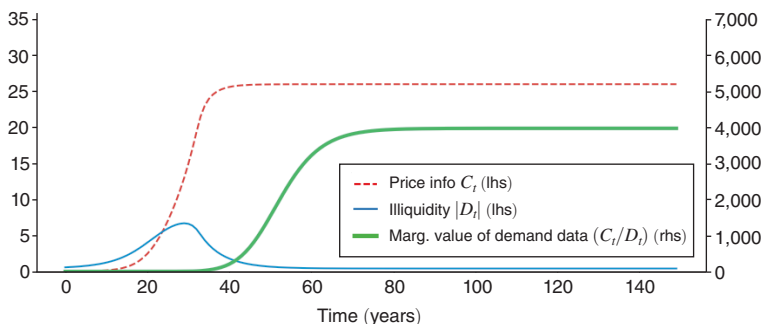


FIGURE 2. PRICE INFORMATION SENSITIVITY (C_t) RISES AND THE PRICE IMPACT OF TRADES ($|D_t|$) STAGNATES

Notes: C_t is the impact of future dividend innovations on price. ($|D_t|$) is the price impact of a one-unit uninformed trade. This illiquidity measure is nonmonotonic, despite the rise in market-making (demand-driven) trades. $(C_t/D_t)^2$ is the marginal value of demand information relative to fundamental information. The x -axis is time in years.

the dynamic force $-\rho \text{var}[p_{t+1} + d_{t+1} | \bar{\mathcal{I}}_t]$, which increases it, are nearly canceling each other out.

Price and Return Volatility.—One might think that future information risk implies an unrealistically high future price or return volatility. It does not. Between the start of the simulation and period 18, which corresponds to 2015, the price variance rises by only 3.4 percent. (See online Appendix Section C.7 for the complete volatility time series.) That is less than 0.2 percent annually. Such a minute increase in price volatility is sufficient to offset the elimination of most dividend uncertainty, because prices are so much larger in magnitude than dividends.

In the return space, the change is even less detectable. When we examine the volatility of the returns, we find no significant time trend in either the model or the data. To substantiate this claim, we simulate prices from the model, inflate them with a consumer price index, and then construct a return series: $(p_{t+1} + d_{t+1})/p_t$. We compare this model return to an empirical return series, derived from the monthly S&P 500 price index (1980–2015). For both the model and data, we estimate a GARCH model to calculate volatility for each month.¹⁶ When we regress this GARCH-implied return variance σ_t on time (monthly), the average variance is 0.005, but the coefficient on time is zero out to 5 digits. Because of the large number of observations that we can simulate, this number is statistically different from zero. But its economic magnitude is trivial. Variance of the S&P 500 returns has a time trend of 0.00001. That coefficient is statistically indistinguishable from zero. In short, the return variance, in both the model and the data, appears stable over time.

¹⁶The dividends are imputed, as described in the calibration section. The equation we estimate is a TGARCH(1,1), which takes the form, $\sigma_{\eta|t-1}^2 = \omega + (\alpha + \gamma \cdot \mathbf{1}_{(r_{t-1} < 0)})r_{t-1}^2 + \beta\sigma_{\eta|t-2}^2$. It allows for positive and negative returns to affect volatility differently. We estimate the coefficients by maximum likelihood.

E. Which Trends Come from Where?

The rise in informativeness and the stagnation of liquidity come from the combination of changes in fundamental and demand data, both today and in the future. To understand the source of these results, we turn off the growth in one of these types of data and see how the resulting price coefficients differ.¹⁷ Online Appendix Section D reports the full set of results from each exercise, with figures.

In a first exercise, designed to isolate the effect of demand data, we set demand analysis to zero and keep fundamental analysis on the same trajectory that it follows in the full model. We also do a version of this exercise where all of the growing data capacity is allocated to fundamental analysis ($\Omega_f = \sqrt{K_t}$). Given these exogenous information sequences, we resolve for the equilibrium pricing coefficients in the equilibrium of the financial market. We learn that demand data make prices only slightly more sensitive to new information. The effect on the price sensitivity to information, C_t , is negligible throughout most of the path. Demand data's effect is small because by the time demand data become prevalent, C_t is nearly at its maximum level. In contrast, demand data govern market liquidity. Without it, the price impact at the end of the simulation would be twice as large as the end of the path, when demand data are abundant. Moreover, we do see demand data affect forecast accuracy, defined as the conditional precision of the future dividend. By the end of the path, nearly all the accuracy of the dividend forecasts is due to demand data.

In a second exercise, we do the reverse: we fix fundamental data at a low level ($\Omega_{ft} = 0.01$) and keep demand data on its previous equilibrium trajectory. We cannot set it to 0, because otherwise $C_t = 0$, mechanically for any level of demand analysis. When information is scarce and fundamental analysis is dominant, removing fundamental data significantly hurts price sensitivity to dividends (C_t), but it has little effect on liquidity or informativeness in the long run. As technology improves, demand analysis becomes dominant. Thus, while some fundamental information is essential, its growth has little quantitative effect on prices in the long run.

A third exercise isolates the effect of long-lived assets and risk from future prices. The most important role of long-lived assets is that they preserve equilibrium existence. When we switch to computing the static model by setting $\pi = 0$, we can only simulate the first few periods before the static solution ceases to exist. The price coefficients have similar dynamics. The key difference is that the dynamic model price is more sensitive to dividend innovations. The reason is that for a long-lived asset, dividends signal not only a rise in current dividends, but also a rise in expected future dividends. That has a larger effect on the expected asset value and thus on price.

F. What Is Robust? What Is Fragile?

These numerical results are simply examples. Some of the features they illustrate are robust to other parameter values, others are not. The online Appendix explores results with different risk aversions, variances of dividends and demand

¹⁷We thank our anonymous referees for suggesting these exercises.

shocks, and rates of time preference. What is consistent is that demand analysis always rises, price informativeness always rises, and the marginal value of demand information $(C/D)^2$ also always rises. Quantitatively, Ω_{xt} consistently surpasses Ω_{ft} once C_t/D_t crosses $\sqrt{\chi_x}$. (See online Appendix Section E for details).

Unbalanced Technological Change.—We chose to model technological progress in a way that increases the potential precision of fundamental or demand information equally. We made this choice because otherwise we would not know which results come from the technological process and which arise from its imbalanced nature. But it is quite possible that technological progress has not been balanced. To explore this possibility, we consider a world where demand data processing efficiency grows more quickly.

When demand analysis efficiency is low, investors analyze fundamental data, just like before. The reason is that no matter how efficient, demand data processing still has zero marginal value at zero data. But as demand analysis becomes relatively cheaper, fundamental analysis falls by more than before. In the long run, the marginal benefit of both types of analysis converges to a constant ratio, as before. (See online Appendix Figure 18.)

In short, most of our conclusions are unaltered. Liquidity is still flat. However, if we see stagnating market efficiency, this is consistent with a world where demand analysis efficiency is improving at a faster rate and is displacing fundamental analysis.

IV. Suggestive Evidence

The shift from fundamental to demand analysis in our model should appear empirically as a change in investment strategies. Indeed, there is some evidence that funds have shifted their strategies in a way that is consistent with our predictions. In the TASS database, many hedge funds report that their fund has a “fundamental,” “mixture,” or “quantitative” strategy. Since 2000, the assets under management of fundamental funds, whether measured by fund or in total, has waned (Figure 3). Instead, other strategies, some based on market data, are surging.¹⁸

A different indicator that points to the growing importance of demand data comes from the frequency of web searches. From 2004 to 2016, the frequency of Google searches for information about “order flow” has risen roughly three-fold (online Appendix Figure 20). This is not an overall increase in attention to asset market information. In contrast, the frequency of searches for information about “fundamental analysis” fell by about one-half over the same time period.

In practice, much of the trade against order flow takes the form of algorithmic trading, for a couple of reasons. First, while firm fundamentals are slow-moving, demand can reverse rapidly. Therefore, mechanisms that allow traders to trade

¹⁸Source: Lipper TASS. The data are monthly from 1994–2015. The database reports on 17,534 live and defunct funds. Quantitative or “quant” strategies are defined in Investopedia as follows: “Quantitative trading techniques include high-frequency trading, algorithmic trading and statistical arbitrage. These techniques are rapid-fire and typically have short-term investment horizons. Many quantitative traders are more familiar with quantitative tools, such as moving averages and oscillators.” The Lipper TASS database is from Thomson Reuters (1994–2015) and can be accessed via Wharton Research Data Services (WRDS).




FIGURE 3. HEDGE FUNDS ARE SHIFTING AWAY FROM FUNDAMENTAL ANALYSIS

Notes: Database reports on 17,534 live and defunct funds. The Lipper TASS database is from Thomson Reuters (1994–2015) and can be accessed via Wharton Research Data Services (WRDS).

Source: Lipper TASS. Data are monthly assets under management per fund, from 1994–2015.

quickly are more valuable for fast-moving, demand-based strategies. Second, while fundamental information is more likely to be textual, partly qualitative, and varied in nature, demand is more consistently data-oriented and therefore more amenable to algorithmic analysis. Hendershott, Jones, and Menkveld (2011) measures algorithmic trading and finds that it has increased, but this increase occurred most rapidly between the start of 2001 and the end of 2005. Over this period, average trade size fell and algorithmic trading increased about seven-fold, consistent with model's predictions for demand-based trading strategies.

Asset price trends are also consistent with our predictions. Bai, Philippon, and Savov (2016) measures a long-run rise in equity price informativeness. They measure price informativeness using a coefficient from a regression of future earnings (at the 1-year, 3-year, and 5-year horizons) on the current ratio of market value to book value. Over the 1960–2010 period, they find a 60 percent rise in the three-year price informativeness and an 80 percent rise in the five-year price informativeness, both of which are highly statistically significant. Similarly, studying liquidity over the last century, Jones (2002) finds a great deal of cyclical variation, but little trend in liquidity, as measured by bid-ask spreads.¹⁹

Our claim is not that our model is the primary explanation for these phenomena, or that we can match the timing or magnitude of the increases or decreases. We only wish to suggest that our predictions are not obviously at odds with long-run financial trends.

V. Conclusion

We explore the consequences of a simple deterministic increase in the productivity of information processing in the financial sector. We find that when the financial sector becomes more efficient at processing information, it changes the incentives

¹⁹Recent work by Kojien and Yogo (2016) measures a large fall in the price impact of institutional traders. This may not be inconsistent with our results, for two reasons. First, our liquidity measure is the price impact of a non-informational trade. That is not the same as the price impact of an institutional trader, who will often be trading on information. Second, in many cases, institutional traders have reduced their price impact by finding uninformed demand to trade against. To the extent that the reduced price impact reflects more market making and less direct trading on information, this reduced impact is consistent with our long-run demand analysis trend.

to acquire information about future dividends (fundamentals) versus demand (nonfundamental shocks to price). Thus, a simple rise in information processing productivity can explain the transformation of financial analysis from a sector that primarily investigates the fundamental profitability of firms to a sector that does a little fundamental analysis but that mostly concentrates on acquiring and processing client demand. This is consistent with suggestive evidence that the nature of financial analysis and its associated trading strategies have changed.

Many feared that this technological transformation would harm market efficiency, while others argued that markets are more liquid/efficient than ever before. Neither phenomenon is a logical consequence of information technology. Although fundamental analysis declines, price informativeness still rises. The reason is that even if many traders are extracting others' information, this still makes the average trader better informed and the price more informative. But the benefits of this technological transformation may also be overstated. The promise that traders would be standing ready to take the other side of uninformed trades would improve market liquidity is only half of the story. What this narrative misses is that more informed traders in the future make prices react more strongly to new information, which makes future asset values riskier. This increase in risk makes traders move market prices by more and pushes market liquidity back down. The net effect could go either way.

Of course, there are many other features one might want to add to this model to speak to other related trends in financial markets. One might make fundamental changes more persistent than demand innovations so that different styles of trade are associated with different trading volumes. Another possibility is to explore regions in this model where the equilibrium does not exist and use its nonexistence as the basis for a theory of market breakdowns or freezes. Another extension might ask where demand signals come from. In practice, people observe demand data because they intermediate trades. Thus, the value of demand information might form the basis for a new theory of intermediation. In such a world, more trading might well generate more information for intermediaries and faster or stronger responses by market participants to changes in market conditions. Finally, one might regard this theory as a prescriptive theory of optimal investment, compare it to investment practice, and compute expected losses from suboptimal information and portfolio choices. For example, a common practice now is to blend fundamental and demand trading by first selecting good fundamental investment opportunities and then using demand information to time the trade. One could construct such a strategy in this model, compare it to the optimal blend of trading strategies, and see if the optimal strategy performs better on market data.

While this project, with its one simple driving force, leaves many questions unanswered, it also provides a tractable foundation on which to build, to continue exploring how and why asset markets are evolving as financial technology improves.

APPENDIX A: MODEL SOLUTION DETAILS

A. Bayesian Updating

To form the conditional expectation, $E[f_{it}|\mathcal{I}_{it}]$, we need to use Bayes' law. But first, we need to know what signal investors extract from price, given their

demand signal, η_x . We can rearrange linear price equation (8) to write a function of the price as the dividend innovation plus the mean zero noise: $\eta_{pit} = y_{t+1} + (D_t/C_t)(x_{t+1} - E[x_{t+1} | \eta_{xit}])$, where the price signal and the signal precision are

$$(28) \quad \eta_{pit} \equiv (p_t - A_t - B(d_t - \mu) - D_t E[x | \eta_{xit}]) / C_t,$$

$$(29) \quad \Omega_{pit} \equiv (C_t/D_t)^2 (\tau_x + \Omega_{xit}).$$

For the simple case of an investor who learned nothing about demand ($E[x] = 0$), the information contained in prices is $(p_t - A_t - B(d_t - \mu)) / C_t$, which is equal to $y_{t+1} + D_t/C_t x_{t+1}$. Since x_{t+1} is a mean-zero random variable, this is an unbiased signal of the asset dividend innovation, y_{t+1} . The variance of the signal noise is $\text{var}[D/Cx] = (D/C)^2 \tau_x^{-1}$. The price signal precision Ω_{pit} is the inverse of this variance.

But conditional on η_{xit} , x_{t+1} is typically not a mean-zero random variable. Instead, investors use Bayes' law to combine their prior, that $x_{t+1} = 0$, with precision τ_x with their demand signals: η_{xit} with precision Ω_{xit} . The posterior mean and variance are

$$(30) \quad E[x | \eta_{xit}] = \frac{\Omega_{xit} \eta_{xit}}{\tau_x + \Omega_{xit}},$$

$$(31) \quad V[x | \eta_{xit}] = \frac{1}{\tau_x + \Omega_{xit}}.$$

Since that is equal to $y_{t+1} + D_t/C_t(x_{t+1} - E[x_{t+1} | \eta_{xit}])$, the variance of price signal noise is $(D_t/C_t)^2 \text{var}[x_{t+1} | \eta_{xit}]$. In other words, the precision of the price signal for agent i (and therefore for every agent, since we are looking at symmetric information choice equilibria) is $\Omega_{pit} \equiv (C_t/D_t)^2 (\tau_x + \Omega_{xit})$.

Now, we can again use Bayes' law for normal variables to form beliefs about the asset payoff. We combine the prior μ , price/demand information η_{pit} , and fundamental signal η_{fit} into a posterior mean and variance:

$$(32) \quad E[y_{t+1} | \mathcal{I}_{it}] = (\tau_0 + \Omega_{pit} + \Omega_{fit})^{-1} (\Omega_{pit} \eta_{pit} + \Omega_{fit} \eta_{fit}),$$

$$(33) \quad V[y_{t+1} | \mathcal{I}_{it}] = (\tau_0 + \Omega_{pit} + \Omega_{fit})^{-1}.$$

Average Expectations and Precisions: Next, we integrate over investors i to get the average conditional expectations. Begin by considering the average price information. The price information content is $\Omega_{pit} \equiv (C_t/D_t)^2 (\tau_x + \Omega_{xit})$. In principle, this can vary across investors. But since all are ex ante identical, they make identical information decisions. Thus, $\Omega_{pit} = \Omega_{pt}$ for all investors i . Since this precision is identical for all investors, we drop the i subscript in what follows. But the realized price signal still differs because signal realizations are heterogeneous. Since the signal precisions are the same for all agents, we can just integrate over signals to get the average signal: $\int \eta_{pit} di = (1/C_t)(p_t - A_t - B(d_t - \mu)) - (D_t/C_t) \times \text{var}(x_{t+1} | \bar{\mathcal{I}}_t) \Omega_{xit} x_{t+1}$. Since $\Omega_{pit}^{-1} = (D_t/C_t)^2 \text{var}(x_{t+1} | \bar{\mathcal{I}}_t)$, we can rewrite this as

$$(34) \quad \int \eta_{pit} di = \frac{1}{C_t} (p_t - A_t - B(d_t - \mu)) - \frac{C_t}{D_t} \Omega_{pit}^{-1} \Omega_{xit} x_{t+1}.$$

Next, let's define some conditional variance/precision terms to simplify notation. The first term, Ω_t , is the precision of the future price plus dividend (the asset payoff). It comes from taking variance of pricing equation (8). It turns out that variance Ω_t^{-1} can be decomposed into a sum of two terms. The first, \hat{V}_t , is the variance of the dividend innovation. This variance depends on information choices Ω_{ft} and Ω_{xt} . The other term, Z_t , depends on future information choices through the $t + 1$ price coefficients. Formally,

$$(35) \quad \hat{V}_t \equiv \text{var}(y_{t+1}|\bar{\mathcal{I}}_t) = (\tau_0 + \Omega_{ft} + \Omega_{pt})^{-1} = (\tau_0 + \Omega_{ft} + (C_t/D_t)^2(\tau_x + \Omega_{xt}))^{-1},$$

$$(36) \quad \Omega_t^{-1} \equiv \text{var}[\pi p_{t+1} + d_{t+1}|\bar{\mathcal{I}}_t] = \pi C_{t+1}^2 \tau_0^{-1} + \pi D_{t+1}^2 \tau_x^{-1} + (1 + \pi B_{t+1})^2 \hat{V}_t,$$

$$(37) \quad Z_t = \frac{\pi \rho}{r} (r - \pi G) (C_{t+1}^2 \tau_0^{-1} + D_{t+1}^2 \tau_x^{-1}),$$

$$(38) \quad \Omega_t^{-1} = \frac{r}{\rho(r - \pi G)} Z_t + \left(\frac{r}{r - \pi G}\right)^2 \hat{V}_t.$$

Thus, $Z_t = 0$ if $\pi = 0$.

The last equation, (38), shows the relationship between Ω , \hat{V} , and Z . This decomposition is helpful because we will repeatedly take derivatives where we take future choices (Z_t) as given and vary current information choices (\hat{V}).

Next, we can compute the average expectations

$$(39) \quad \int E[y_{t+1}|\bar{\mathcal{I}}_t] di = \hat{V}_t \left[\Omega_{ft} y_{t+1} + \Omega_{pt} \left(\frac{1}{C_t} (p_t - A_t - B(d_t - \mu)) - \frac{C_t}{D_t} \Omega_{pt}^{-1} \Omega_{xt} x_{t+1} \right) \right]$$

$$(40) \quad = \hat{V}_t \left[\Omega_{ft} y_{t+1} + \Omega_{pt} \frac{1}{C_t} (p_t - A_t - B(d_t - \mu)) - \frac{C_t}{D_t} \Omega_{xt} x_{t+1} \right],$$

$$(41) \quad \int E[\pi p_{t+1} + d_{t+1}|\bar{\mathcal{I}}_t] di = A_{t+1} + (1 + \pi B) E[d_{t+1}|\bar{\mathcal{I}}_t] \\ = A_{t+1} + (1 + \pi B) (\mu + G(d_t - \mu) + E[y_{t+1}|\bar{\mathcal{I}}_t]).$$

B. Solving for Equilibrium Prices

The price conjecture is

$$(42) \quad p_t = A_t + B_t(d_t - \mu) + C_t y_{t+1} + D_t x_{t+1}.$$

We will solve for the prices for the general supply of asset, \bar{x} , although in the main text it is normalized to one unit.

The price coefficients solve the system of recursive equations:

$$(43) \quad A_t = \frac{1}{r} \left[\pi A_{t+1} + \mu - \rho \bar{x} \left(\pi C_{t+1}^2 \tau_0^{-1} + \pi D_{t+1}^2 \tau_x^{-1} + (1 + \pi B_{t+1})^2 \left(\tau_0 + \left(\frac{K}{1 + \frac{1}{\chi_x} \xi^4} \right)^{\frac{1}{2}} + \xi^2 \left(\tau_x + \frac{\xi^2}{\chi_x} \left(\frac{K}{1 + \frac{1}{\chi_x} \xi^4} \right)^{\frac{1}{2}} \right) \right)^{-1} \right) \right],$$

$$(44) \quad B_t = \frac{1}{r} (1 + \pi B_{t+1}) G = \frac{G}{r - \pi G},$$

$$(45) \quad C_t = \frac{1}{r - \pi G} \left(1 - \tau_0 \left(\tau_0 + \left(\frac{K}{1 + \frac{1}{\chi_x} \xi^4} \right)^{\frac{1}{2}} + \xi^2 \left(\tau_x + \frac{\xi^2}{\chi_x} \left(\frac{K}{1 + \frac{1}{\chi_x} \xi^4} \right)^{\frac{1}{2}} \right) \right)^{-1} \right),$$

$$(46) \quad D_t = \left(\frac{\tau_x}{r - \pi G} \xi - \frac{\rho r}{(r - \pi G)^2} \right) \times \left(\tau_0 + \left(\frac{K}{1 + \frac{1}{\chi_x} \xi^4} \right)^{\frac{1}{2}} + \xi^4 \left(\tau_x + \frac{\xi^2}{\chi_x} \left(\frac{K}{1 + \frac{1}{\chi_x} \xi^4} \right)^{\frac{1}{2}} \right) \right)^{-1} - \pi \frac{\rho}{r} (C_{t+1}^2 \tau_0^{-1} + D_{t+1}^2 \tau_x^{-1}),$$

where $\xi = C_t/D_t$ denotes the date- t signal to noise ratio, which is the solution to equation (69) in the online Appendix. The high- K limit pricing coefficients are the fixed point of the system above.

The sequence of pricing coefficients is known at every date. Signals η_{fit} and η_{xit} are the same as before, except that their precisions, Ω_{ft} and Ω_{xt} , may change over time if that is the solution to the information choice problem.

The conditional expectation and variance of y_{t+1} (32) and (33) are the same, except that the Ω_{pt} term gets a t subscript now because $\Omega_{pt} \equiv (C_t/D_t)^2 (\tau_x + \Omega_{xt})$. Likewise, the mean and the variance of x_{t+1} (30) and (31) are the same with a time-subscripted Ω_{xt} . Thus, the average signals are the same with t -subscripts:

$$(47) \quad \int \eta_{pit} di = \frac{1}{C_t} (p_t - A_t - B_t(d_t - \mu)) - \frac{D_t}{C_t} \text{var}(x_{t+1} | \bar{I}_t) \Omega_{xt} x_{t+1}.$$

Since $\Omega_{pt}^{-1} = (D_t/C_t)^2 \text{var}(x_{t+1} | \bar{I}_t)$, we can rewrite this as

$$(48) \quad \int \eta_{pit} di = \frac{1}{C_t} (p_t - A_t - B_t(d_t - \mu)) - \frac{C_t}{D_t} \Omega_{pt}^{-1} \Omega_{xt} x_{t+1}.$$

Solving for Nonstationary Equilibrium Prices.—To solve for equilibrium prices, start from the portfolio first-order condition for investors (7) and equate total demand with total supply. The total risky asset demand (excluding noisy demand) is

$$(49) \quad \int q_{it} di = \frac{1}{\rho} \Omega_t \left[\pi A_{t+1} + (1 + \pi B_{t+1}) (\mu + G(d_t - \mu)) + \hat{V}_t \left[\Omega_{ft} y_{t+1} + \Omega_{pt} \frac{1}{C_t} (p_t - A_t - B_t(d_t - \mu)) - \frac{C_t}{D_t} \Omega_{xt} x_{t+1} \right] - \pi B_{t+1} \mu - p_t r \right].$$

The market clearing condition equates the expression above to the residual asset supply $\bar{x} + x_{t+1}$. The model assumes the asset supply is 1. We use the notation \bar{x} here for more generality because then we can apply the result to the model with issuance costs where asset supply is a choice variable. Rearranging the market clearing condition (just multiplying through by $\rho \Omega_t^{-1}$ and bringing the p terms to the left) yields

$$(50) \quad \left[r - (1 + \pi B_{t+1}) \hat{V}_t \Omega_{pt} \frac{1}{C_t} \right] p_t = -\rho \Omega_t^{-1} (\bar{x} + x_{t+1}) + \pi A_{t+1} + (1 + \pi B_{t+1}) (\mu + G(d_t - \mu)) + (1 + \pi B_{t+1}) \hat{V}_t \Omega_{ft} y_{t+1} - (1 + \pi B_{t+1}) \hat{V}_t \Omega_{pt} \frac{1}{C_t} (A_t + B_t(d_t - \mu)) - (1 + \pi B_{t+1}) \frac{C_t}{D_t} \hat{V}_t \Omega_{xt} x_{t+1} - \pi B_{t+1} \mu.$$

Solve for p to get

$$\begin{aligned} & \left(r - (1 + \pi B_{t+1}) \hat{V}_t \Omega_{pt} \frac{1}{C_t} \right) (A_t + B_t(d_t - \mu) + C_t y_{t+1} + D_t x_{t+1}) \\ & = -\rho \Omega_t^{-1} (\bar{x} + x_{t+1}) + \pi A_{t+1} + (1 + \pi B_{t+1}) (\mu + G(d_t - \mu)) \\ & \quad + (1 + \pi B_{t+1}) \hat{V}_t \Omega_{ft} y_{t+1} - (1 + \pi B_{t+1}) \hat{V}_t \Omega_{pt} \frac{1}{C_t} (A_t + B_t(d_t - \mu)) \\ & \quad - (1 + \pi B_{t+1}) \frac{C_t}{D_t} \hat{V}_t \Omega_{xt} x_{t+1} - \pi B_{t+1} \mu. \end{aligned}$$

Multiply both sides by the first term on the left-hand side and match the coefficients to get

$$(51) \quad A_t = \left[r - (1 + \pi B_{t+1}) \hat{V}_t \Omega_{pt} \frac{1}{C_t} \right]^{-1} \left(-\rho \Omega_t^{-1} \bar{x} + \pi A_{t+1} + (1 + \pi B_{t+1}) \mu - (1 + \pi B_{t+1}) \hat{V}_t \Omega_{pt} \frac{1}{C_t} A_t - \pi B_{t+1} \mu \right).$$

Multiplying both sides by the inverse term:

$$\begin{aligned} r A_t - (1 + \pi B_{t+1}) \hat{V}_t \Omega_{pt} \frac{1}{C_t} A_t & = -\rho \Omega_t^{-1} \bar{x} + \pi A_{t+1} + (1 + \pi B_{t+1}) \mu \\ & \quad - (1 + \pi B_{t+1}) \hat{V}_t \Omega_{pt} \frac{1}{C_t} A_t - \pi B_{t+1} \mu \end{aligned}$$

and canceling the $(1 + \pi B_{t+1}) \hat{V}_t \Omega_{pt} (1/C_t) A_t$ term on both sides leaves

$$(52) \quad \begin{aligned} rA_t &= -\rho \Omega_t^{-1} \bar{x} + \pi A_{t+1} + (1 + \pi B_{t+1})\mu - \pi B_{t+1}\mu, \\ A_t &= \frac{1}{r} [\pi(A_{t+1} - B_{t+1}\mu) + (1 + \pi B_{t+1})\mu - \rho \Omega_t^{-1} \bar{x}] \\ &= \frac{1}{r} [\pi A_{t+1} + \mu - \rho \Omega_t^{-1} \bar{x}]. \end{aligned}$$

Risk Premium.—The risk premium is defined as

$$(53) \quad rp_t = \frac{E[p_{t+1} + d_{t+1}]}{E[p_t]} - r.$$

The risk premium can be written as

$$rp_t = \frac{A_{t+1} + \mu}{A_t} - r = \frac{r(A_{t+1} + \mu)}{A_{t+1} + \mu - \rho \Omega_t^{-1} \bar{x}} - r = \frac{r\rho \bar{x} \Omega_t^{-1}}{A_{t+1} + \mu - \rho \Omega_t^{-1} \bar{x}}$$

where the first equality takes the unconditional expectation, recognizing that $E[d_t] = \mu$, and the second equation uses the derivation of A_t in equation (52). Note that if all the variance goes to zero, $\Omega_t^{-1} \rightarrow 0$, the risk premium also goes to zero.

Note that because, $\bar{x} = 1$ for the main model, in the main text, in equation (9) \bar{x} is set to 1.

Matching coefficients on d_t yields

$$(54) \quad B_t = \left[r - (1 + \pi B_{t+1}) \hat{V}_t \Omega_{pt} \frac{1}{C_t} \right]^{-1} \left[(1 + \pi B_{t+1})G - (1 + \pi B_{t+1}) \hat{V}_t \Omega_{pt} \frac{B_t}{C_t} \right].$$

Multiplying on both sides by the inverse term

$$(55) \quad rB_t - (1 + \pi B_{t+1}) \hat{V}_t \Omega_{pt} \frac{1}{C_t} B_t = (1 + \pi B_{t+1})G - (1 + \pi B_{t+1}) \hat{V}_t \Omega_{pt} \frac{B_t}{C_t}$$

and canceling the last term on both sides yields

$$(56) \quad B_t = \frac{1}{r} (1 + \pi B_{t+1})G.$$

As long as r and G do not vary over time, a stationary solution for B exists. That stationary solution would be (10).

Next, collecting all the terms in y_{t+1} ,

$$(57) \quad C_t = \left[r - (1 + \pi B_{t+1}) \hat{V}_t \Omega_{pt} \frac{1}{C_t} \right]^{-1} (1 + \pi B_{t+1}) \hat{V}_t \Omega_{ft},$$

multiplying both sides by the first term inverse yields $rC_t - (1 + \pi B_{t+1}) \hat{V}_t \Omega_{pt} = (1 + \pi B_{t+1}) \hat{V}_t \Omega_{ft}$. Then, dividing through by r and collecting terms in $\hat{V}(1 + \pi B_{t+1})$ yields $C_t = (1/r)(1 + \pi B_{t+1}) \hat{V}_t (\Omega_{pt} + \Omega_{ft})$. Next, using the fact that $\hat{V}^{-1} = \tau_0 + \Omega_{pt} + \Omega_{ft}$, we get $C_t = 1/r(1 + \pi B_{t+1})(1 - \tau_0 \hat{V}_t)$. Of course

the \hat{V}_t term has C_t and D_t in it. If we use the stationary solution for B (if r and G do not vary), then we can simplify to get

$$(58) \quad C_t = \frac{1}{r - \pi G} (1 - \tau_0 \hat{V}_t).$$

Finally, we collect terms in x_{t+1} ,

$$(59) \quad D_t = \left[r - (1 + \pi B_{t+1}) \hat{V}_t \Omega_{pt} \frac{1}{C_t} \right]^{-1} \left[-\rho \Omega_t^{-1} - (1 + \pi B_{t+1}) \frac{C_t}{D_t} \hat{V}_t \Omega_{xt} \right],$$

multiply by the inverse term, and then use $\Omega_{pt} = (C_t/D_t)^2 (\tau_x + \Omega_{xt})$ to get

$$(60) \quad r D_t - (1 + \pi B_{t+1}) \hat{V}_t \frac{C_t}{D_t} (\tau_x + \Omega_{xt}) = -\rho \Omega_t^{-1} - (1 + \pi B_{t+1}) \frac{C_t}{D_t} \hat{V}_t \Omega_{xt}.$$

Then, adding $(1 + B)(C_t/D_t) \hat{V}_t \Omega_{xt}$ to both sides, and substituting in B (stationary solution), we get

$$(61) \quad \begin{aligned} D_t &= \frac{1}{r - \pi G} \hat{V}_t \tau_x \frac{C_t}{D_t} - \frac{\rho}{r} \Omega_t^{-1}, \\ D_t &= \frac{1}{r - \pi G} \left[\left(\tau_x \frac{C_t}{D_t} - \frac{r\rho}{r - \pi G} \right) \hat{V}_t - Z_t \right]. \end{aligned}$$

Of course, D_t still shows up quadratically, and also in \hat{V}_t . The future coefficient values C_{t+1} and D_{t+1} show up in Ω_t .

C. Solving Information Choices

Details of Step 3: Compute ex ante expected utility. Note that the expected excess return $(E[\pi p_{t+1} + d_{t+1} | \mathcal{I}_{it}] - p_t r)$ depends on fundamental and supply signals, and on prices, all of which are unknown at the beginning of the period. Because asset prices are linear functions of normally distributed shocks, $E[\pi p_{t+1} + d_{t+1} | \mathcal{I}_{it}] - p_t r$ is normally distributed as well.

With $E \ln E$ preferences, $(E[\pi p_{t+1} + d_{t+1} | \mathcal{I}_{it}] - p_t r) \Omega (E[\pi p_{t+1} + d_{t+1} | \mathcal{I}_{it}] - p_t r)$ is a non-central χ_x^2 -distributed variable. Computing its mean yields a first term that depends on known endowment e_{it} and on terms that depend on information: $\rho e_{it} + \rho E[q_{it}(E[\pi p_{t+1} + d_{t+1} | \mathcal{I}_{it}] - p_t r) | \mathbf{I}_t^-] - (\rho^2/2) E[q_{it}^2 \text{var}[\pi p_{t+1} + d_{t+1} | \bar{\mathcal{I}}_t]^{-1} | \mathbf{I}_t^-]$. As we argue in the main text, $\text{var}[\pi p_{t+1} + d_{t+1} | \bar{\mathcal{I}}_t]$ depends only on the posterior variance of the total payoff, Ω^{-1} . Note, Ω_{ft} and Ω_{ft} do not enter separately.

With the expected utility, $(E[\pi p_{t+1} + d_{t+1} | \mathcal{I}_{it}] - p_t r) \Omega (E[\pi p_{t+1} + d_{t+1} | \mathcal{I}_{it}] - p_t r)$ is still a non-central χ_x^2 -distributed variable. But the expected utility is the expectation of the exponential of this expression: $E[\exp\{(E[\pi p_{t+1} + d_{t+1} | \mathcal{I}_{it}] - p_t r) \times \Omega (E[\pi p_{t+1} + d_{t+1} | \mathcal{I}_{it}] - p_t r)\} | \mathbf{I}_t^-]$. The exponential of a chi-square distribution is a Wishart. Expected utility is the mean of this expression.

The end-of-period expectation of excess return is distributed $(E[\pi p_{t+1} + d_{t+1} | \mathcal{I}_{it}] - p_t r) \sim N((1 - C)\mu - A, V_{ER})$ where V_{ER} , as in the $E \ln E$ model, is an increasing function of payoff precision Ω , and does not contain terms

in Ω_{ft} or Ω_{xt} , except through Ω . Using the formula for the mean of a Wishart (see Veldkamp 2011, Appendix Ch.7), we compute the ex ante expected utility:

$$(62) \quad U = \frac{1}{2} \text{tr}(\Omega V_{ER}) + \frac{1}{2} ((1 - C)\mu - A)' \Omega ((1 - C)\mu - A).$$

Since the precisions Ω_{ft} and Ω_{xt} only enter expected the utility through the posterior precision of payoffs Ω , the same is true for the exponential of this expression. Since the exponential function is a monotonic increasing function, we know that expected utility takes the form of an increasing function of Ω . As long as Ω is a sufficient statistic for the data choices in utility, investors' data choices that maximize Ω also maximize the expected utility.

Details of Step 4: Solve for fundamental information choices. Note that in expected utility (15), choice variables Ω_{ft} and Ω_{xt} enter only through posterior variance Ω^{-1} and through $V[E[\pi p_{t+1} + d_{t+1} | \mathcal{I}_{it}] - p_t r | \mathbf{I}_t^-] = V[\pi p_{t+1} + d_{t+1} - p_t r | \mathbf{I}_t^-] - \Omega_t^{-1}$. Since there is a continuum of investors, and since $V[\pi p_{t+1} + d_{t+1} - p_t r | \mathbf{I}_t^-]$ and $E[E[\pi p_{t+1} + d_{t+1} | \mathcal{I}_{it}] - p_t r | \mathbf{I}_t^-]$ depend only on $t - 1$ variables and parameters, and on aggregate information choices, each investor takes them as given. If the objective is to maximize an increasing function of Ω , then the information choices must maximize Ω as well.

REFERENCES

- Abis, Simona. 2017. "Market Timing in Bayesian Portfolio Optimization." Unpublished.
- Amador, Manuel, and Pierre-Olivier Weill. 2010. "Learning from Prices: Public Communication and Welfare." *Journal of Political Economy* 118 (5): 866–907.
- Angeletos, George-Marios, and Jennifer La'O. 2013. "Sentiments." *Econometrica* 81 (2): 739–79.
- Asriyan, Vladimir, and Victoria Vanasco. 2014. "Informed Intermediation over the Cycle." Unpublished.
- Atkeson, Andrew, and Robert E. Lucas Jr. 1992. "On Efficient Distribution with Private Information." *Review of Economic Studies* 59 (3): 427–53.
- Babus, Ana, and Cecilia Parlatore. 2015. "Strategic Fragmented Markets." Unpublished.
- Bai, Jennie, Thomas Philippon, and Alexi Savov. 2016. "Have Financial Markets Become More Informative?" *Journal of Financial Economics* 122 (3): 625–54.
- Banerjee, Snehal, and Brett Green. 2015. "Signal or Noise? Uncertainty and Learning about Whether Other Traders Are Informed." *Journal of Financial Economics* 117 (2): 398–423.
- Barlevy, Gadi, and Pietro Veronesi. 2000. "Information Acquisition in Financial Markets." *Review of Economic Studies* 67 (1): 79–90.
- Bassetto, Marco, and Carlo Galli. 2017. "Is Inflation Default? The Role of Information in Debt Crises." Unpublished.
- Biais, Bruno, Thierry Foucault, and Sophie Moinas. 2015. "Equilibrium Fast Trading." *Journal of Financial Economics* 116 (2): 292–313.
- Blanchard, Olivier, Jean-Paul L'Huilier, and Guido Lorenzoni. 2013. "News, Noise, and Fluctuations: An Empirical Exploration." *American Economic Review* 103 (7): 3045–70.
- Briggs, Joseph, Andrew Caplin, Daniel Martin, and Christopher Tonetti. 2017. "Due Diligence: Job Search with Rationally Inattentive Workers." Unpublished.
- Brunnermeier, Markus K., Michael Sockin, and Wei Xiong. 2017. "China's Model of Managing the Financial System." Unpublished.
- Cai, Zhifeng. 2016. "Dynamic Information Complementarity in Information Acquisition." Unpublished.
- Caplin, Andrew, Mark Dean, and John Leahy. 2018. "Rational Inattention, Optimal Consideration Sets, and Stochastic Choice." *Review of Economic Studies* 86 (3): 1061–94.
- Caplin, Andrew, John Leahy, and Filip Matejka. 2016. "Rational Inattention and Inference from Market Share Data." Unpublished.
- Cespa, Giovanni, and Xavier Vives. 2012. "Dynamic Trading and Asset Prices: Keynes vs. Hayek." *Review of Economic Studies* 79 (2): 539–80.

- Chahrour, Ryan.** 2014. "Public Communication and Information Acquisition." *American Economic Journal: Macroeconomics* 6 (3): 73–101.
- Chien, YiLi, Harold Cole, and Hanno Lustig.** 2016. "Implications of Heterogeneity in Preferences, Beliefs, and Asset Trading Technologies in an Endowment Economy." *Review of Economic Dynamics* 20: 215–39.
- Cisco.** 2018. "Cisco Global Cloud Index: Forecast and Methodology, 2016–2021." White Paper.
- Cover, Thomas M., and Joy A. Thomas.** 1991. *Elements of Information Theory*. 1st ed. New York: John Wiley and Sons.
- Crouzet, Nicolas, Ian Dew-Becker, and Charles G. Nathanson.** 2016. "A Model of Multi-Frequency Trade." Unpublished.
- David, Joel M., Hugo A. Hopenhayn, and Venky Venkateswaran.** 2016. "Information, Misallocation and Aggregate Productivity." *Quarterly Journal of Economics* 131 (2): 943–1005.
- Dávila, Eduardo, and Cecilia Parlatore.** 2016. "Trading Costs and Informational Efficiency." Unpublished.
- Du, Songzi, and Haoxiang Zhu.** 2017. "What Is the Optimal Trading Frequency in Financial Markets?" *Review of Financial Studies* 84 (4): 1606–51.
- Fajgelbaum, Pablo D., Edouard Schaal, and Mathieu Taschereau-Dumouchel.** 2017. "Uncertainty Traps." *Quarterly Journal of Economics* 132 (4): 1641–92.
- Farboodi, Maryam, and Laura Veldkamp.** 2020. "Replication Data for: Long-Run Growth of Financial Data Technology." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/E114984V3>.
- Ganguli, Jayant Vivek, and Liyan Yang.** 2009. "Complementarities, Multiplicity, and Supply Information." *Journal of the European Economic Association* 7 (1): 90–115.
- Glode, Vincent, Richard C. Green, and Richard Lowery.** 2012. "Financial Expertise as an Arms Race." *Journal of Finance* 67 (5): 1723–59.
- Goldstein, Itay, Emre Ozdenoren, and Kathy Yuan.** 2013. "Trading Frenzies and Their Impact on Real Investment." *Journal of Financial Economics* 109 (2): 566–82.
- Grossman, Sanford J., and Joseph E. Stiglitz.** 1980. "On the Impossibility of Informationally Efficient Markets." *American Economic Review* 70 (3): 393–408.
- He, Zhiguo.** 2009. "The Sale of Multiple Assets with Private Information." *Review of Financial Studies* 22 (1): 4787–4820.
- Hendershott, Terrence, Charles M. Jones, and Albert J. Menkveld.** 2011. "Does Algorithmic Trading Improve Liquidity?" *Journal of Finance* 66 (1): 1–33.
- Hennessy, John L., and David A. Patterson.** 2009. *Computer Architecture: A Quantitative Approach*. Amsterdam: Elsevier.
- Jones, Charles M.** 2002. "A Century of Stock Market Liquidity and Trading Costs." Unpublished.
- Kacperczyk, Marcin, Jaromir Nosal, and Luminita Stevens.** 2015. "Investor Sophistication and Capital Income Inequality." Unpublished.
- Koijen, Ralph S. J., and Motohiro Yogo.** 2016. "An Equilibrium Model of Institutional Demand and Asset Prices." Unpublished.
- Lowery, Richard, and Tim Landoigt.** 2016. "Financial Industry Dynamics." Society for Economic Dynamics 2016 Meeting Paper 1248.
- Mackowiak, Bartosz, and Mirko Wiederholt.** 2012. "Business Cycle Dynamics under Rational Inattention." Unpublished.
- Manzano, Carolina, and Xavier Vives.** 2010. "Public and Private Learning from Prices, Strategic Substitutability and Complementarity, and Equilibrium Multiplicity." CEPR Discussion Paper 7949.
- Myatt, David P., and Chris Wallace.** 2012. "Endogenous Information Acquisition in Coordination Games." *Review of Economic Studies* 79 (1): 340–74.
- Nimark, Kristoffer.** 2008. "Monetary Policy with Signal Extraction from the Bond Market." *Journal of Monetary Economics* 55 (8): 1389–1400.
- Standard & Poor's.** 1985–2015. "Compustat Database." Accessed from Wharton Research Data Services (WRDS) on December 1, 2019. <https://wrds-www.wharton.upenn.edu>.
- Thomson Reuters.** 1994–2015. "Lipper TASS Database." Accessed from Wharton Research Data Services (WRDS) on December 1, 2019. <https://wrds-www.wharton.upenn.edu>.
- Veldkamp, Laura L.** 2011. *Information Choices in Macroeconomics and Finance*. Princeton, NJ: Princeton University Press.
- Wang, Jiang.** 1993. "A Model of Intertemporal Asset Prices under Asymmetric Information." *Review of Economic Studies* 60 (2): 249–82.
- Yang, Liyan, and Haoxiang Zhu.** 2016. "Back-Running: Seeking and Hiding Fundamental Information in Order Flows." Rotman School of Management Working Paper 2583915.

This article has been cited by:

1. Charles I. Jones, Christopher Tonetti. 2020. Nonrivalry and the Economics of Data. *American Economic Review* **110**:9, 2819-2858. [[Abstract](#)] [[View PDF article](#)] [[PDF with links](#)]