

4787392-what-hundreds-of-economic-news-events-say-about-belief-overreaction-in-the-stock

SUMMARY: This paper examines the nature and severity of belief distortions in market reactions to hundreds of economic news events. The authors use a combination of structural estimation and algorithmic machine learning to quantify bias. They find that investors systematically overreact to perceptions about multiple fundamental shocks, which often dampens market volatility rather than amplifying it. The study implies that the stock market can underreact to news, even when investors overreact to all shocks. The authors argue that traditional approaches to measuring overreaction are limited and propose a more comprehensive empirical approach to address these gaps in the literature.

METHOD: The authors use a combination of structural estimation and algorithmic machine learning to measure the nature and severity of belief distortions in market reactions to hundreds of economic news events. They analyze a large dataset of news events and investor sentiment to quantify bias and estimate revisions in the representative investor's perceptions about multiple sources of risk.

KEY FINDINGS:

- Investors systematically overreact to perceptions about multiple fundamental shocks.
- Overreaction often dampens rather than amplifies market volatility.
- The stock market can underreact to news, even when investors overreact to all shocks.
- Traditional regression approaches are limited and do not provide a reliable measure of overreaction.

IMPLICATION FOR TRADING: The findings suggest that traditional approaches to measuring overreaction in the stock market may be flawed. Practitioners should consider using a more comprehensive empirical approach to better understand investor behavior and market dynamics. This could lead to more accurate predictions and improved trading strategies. Additionally, the study highlights the importance of considering multiple sources of risk and the potential for underreaction in the stock market, which could have implications for portfolio management and risk assessment.

NBER WORKING PAPER SERIES

WHAT HUNDREDS OF ECONOMIC NEWS EVENTS SAY
ABOUT BELIEF OVERREACTION IN THE STOCK MARKET

Francesco Bianchi
Sydney C. Ludvigson
Sai Ma

Working Paper 32301
<http://www.nber.org/papers/w32301>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2024, Revised June 2026

We are grateful to Nicolò Ceneri, Doruk Gökalp, Do Lee, Steven Zheng, and Bill Shu for excellent research assistance, and to Nicholas Barberis, Francesca Bastianello, Aditya Chaudhry, Nicola Gennaioli, Samuel Hartzmark, Stefan Nagel, Andrei Shleifer, and seminar participants at the NBER Spring 2024 Behavioral Finance meeting, the NBER December 2024 Big Data, Artificial Intelligence, and Financial Economics meeting, the Luiss Finance Conference in May 2024, the UNC Chapel Hill, LSE, Queen Mary University London, Wharton, and University of Texas Austin for helpful comments. Bianchi and Ludvigson received financial support from the National Science Foundation under Grant 2116641. The views expressed are those of the authors and do not necessarily reflect those of the Federal Reserve Board, the Federal Reserve System, or the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Francesco Bianchi, Sydney C. Ludvigson, and Sai Ma. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

What Hundreds of Economic News Events Say About Belief Overreaction in the Stock Market
Francesco Bianchi, Sydney C. Ludvigson, and Sai Ma
NBER Working Paper No. 32301
April 2024, Revised June 2026
JEL No. G1, G12, G4, G41

ABSTRACT

We measure the nature and severity of a variety of belief distortions in market reactions to hundreds of economic news events by synthesizing structural estimation with algorithmic machine learning to quantify bias. We find that investors systematically overreact to perceptions about multiple fundamental shocks, a phenomenon we show often dampens rather than amplifies market volatility via a shock composition effect. Such effects imply that the stock market can underreact to news, even when investors overreact to all shocks.

Francesco Bianchi
Johns Hopkins University
Department of Economics
and CEPR
and also NBER
francesco.bianchi@jhu.edu

Sai Ma
Federal Reserve Board of Governors
sai.ma@frb.gov

Sydney C. Ludvigson
New York University
Department of Economics
and NBER
sydney.ludvigson@nyu.edu

1 Introduction

The pronounced volatility of world equity markets is difficult to reconcile with textbook models in which the price of a stock is the rational expectation of future cash-flow fundamentals, discounted at a constant rate. These theories imply that stock markets should be far more stable than observed, leading a vast literature to explain “excess” stock market volatility with discount rate variation.¹ But recent advancements in the field of behavioral finance point toward a different explanation, namely that investors may exhibit systematic expectational errors (“belief distortions”) that lead them to overreact to news relevant for cash-flow growth. A standard result is that overreaction amplifies market volatility, offering an explanation for observed equity markets that does not rely on variable discount rates.

Documenting evidence of overreaction (or belief distortions more generally) requires both a measure of what investors subjectively expect, and a benchmark for gauging any distortion in subjective growth expectations. The traditional approach to this problem is to use surveys of analysts or investors to measure subjective expectations, and to use in-sample regressions of survey forecast errors on lagged forecast revisions to measure overreaction. Despite valuable insights, the very simplicity and convenience of this approach necessarily leaves several pertinent questions unanswered: First, which real-world events have historically been responsible for market overreactions, and why? The traditional regression approach leaves the precise news events completely unspecified. Second, what are the perceived primitive shocks that investors respond to when they overreact to news? Existing models, such as those in Bordalo, Gennaioli, La Porta and Shleifer (2019), Nagel and Xu (2022), and Bordalo, Gennaioli, La Porta and Shleifer (2024), feature investors who react to unexpected changes in a single fundamental shock to payout or earnings. These papers do not tell us how evidence on belief overreaction might change in a more general setting where multiple primitive shocks are relevant for payout and valuation. Third, how reliable is the traditional regression methodology? Contrary to the traditional approach, dynamic machine learning

¹For textbook treatments of this issue, see Chapters 7 and 8 of Campbell, Lo and MacKinlay (1997), and Chapter 20 of Cochrane (2005).

algorithms designed to quantify the overall magnitude of distortion in beliefs find little evidence that survey forecast errors are related to lagged forecast revisions (Bianchi, Ludvigson and Ma (2022a)), presenting an empirical challenge to the standard inferences drawn from these regressions.

In this paper we revisit the evidence on belief overreaction to news using a more general empirical approach capable of addressing these gaps in the literature. Our objectives are threefold: (i) measuring the stock market's response to specific news events, (ii) estimating revisions in the representative investor's perceptions about multiple sources of risk as a result of those events, and (iii) gauging the quantitative importance of a range of belief distortions in driving the market's reactions to news.

Our approach has four central ingredients. First, we require high-frequency market reactions to specific news events, which we obtain by studying hundreds of macroeconomic data releases, corporate earnings announcements, and Federal Reserve communications. Second, we use a structural asset pricing model to map these empirical news events into revised investor perceptions about multiple primitive shocks that together span cash-flow and discount-rate news. Third, investor beliefs in the structural model must be allowed to depart from rationality. We specify two broad distortions: the first allows for general over- or underreaction due to distorted perceptions about the aggregate economy's laws of motion, implying investors may misattribute one primitive shock to a mixture of others. The second is summarized by a single estimated scalar parameter, ζ , which controls reactions to all perceived shocks. This formulation nests specific belief formation frameworks, including inattention ($\zeta < 0$; e.g., Sims (2003), Gabaix (2019)), diagnostic expectations ($\zeta > 0$; e.g., Bordalo, Gennaioli and Shleifer (2018), Bordalo et al. (2019), and Bordalo et al. (2024)), and rational expectations ($\zeta = 0$). Fourth, we use the dynamic machine learning methodology of Bianchi et al. (2022a) (BLM1) and Bianchi, Lee, Ludvigson and Ma (2025) to construct an objective baseline of efficient expectation formation. Merging this machine learning output with our structural estimation allows us to identify and quantify distortions in the reactions to real-world news events through the lens of the model. A key precursor to our main findings

is that a model with overreaction to all shocks (i.e., $\zeta > 0$) fits the post-war behavior of the stock market best, and with little to no error.

Our main findings can be summarized as follows.

First, we find that while investors systematically *overreact* to all perceived shocks—our baseline estimate has $\hat{\zeta} > 0$ —the stock market frequently *underreacts* to news, resulting in a post-millennial puzzle of “excess stability” rather than excess volatility. This surprising outcome results from a *shock composition effect*: many real-world news events cause investors to revise their perceptions about multiple fundamental shocks simultaneously, in directions that have counteracting but *asymmetric* implications for valuations. Crucially, these asymmetries arise because the overall dynamic distortion is the product of ζ and a vector of fundamental shocks rather than a single shock. Because the elements of this vector possess different volatility and persistence properties, the resulting distortions generated by $\zeta \neq 0$ are inherently shock-specific even though ζ is not. As an example, suppose an event brings predominantly good news about discount rates alongside partially offsetting bad news about cash flows. If the cash flow shock is intrinsically more volatile or more persistent than the discount rate shock, the same $\zeta > 0$ distortion that generates overreaction to both shocks will produce a larger overreaction to cash flow news than it will to discount rate news. Consequently, the market can rise “too little” because the investor’s revised expectations for earnings are more overly pessimistic than her views on discount rates are overly rosy. Such a muted reaction well describes the stock market’s behavior in several major episodes of post-millennial history, most notably the Global Financial Crisis, in which behavioral overreaction was a force for stability rather than volatility. Across the full post-millennial period, this stabilizing force predominates, implying that the observed market was less volatile than a counterfactual rational market. These dynamics underscore a central irony of our findings: while diagnostic expectations are typically analyzed in a univariate framework where they can only deliver excess volatility, a multi-shock model of belief overreaction successfully explains the data not because it amplifies volatility but because it dampens it.

Second, we show that this shock composition effect, and the excess stability phenomenon

it generates, is a direct result of a strict hierarchy of overreactions across different macroeconomic shocks. While investors overreact to all primitive shocks, they do so most strongly to the highly transitory (short-run) and longer-run components of the payout share of output—the most volatile of our estimated fundamental shocks—far outpacing overreactions to shocks that feed into asset prices through economic growth or discount rates.

Third, despite our finding that excess stability predominates over the full sample, these same model estimates imply that the aggregate stock market can still exhibit “excess volatility” around specific real-world news events. This occurs when overreaction to each shock individually amplifies the effects of all shocks combined. In the data, these large, localized market overreactions are typically driven by overreactions to cash flow news and associated with high-frequency jumps in analyst expectations for earnings relative to aggregate output around major news releases. That investors attend strongly to news about cash flows relative to aggregate output is consistent with evidence that the earnings share of output is highly volatile and has contributed more in the long-run to stock market valuations than either economic growth or discount rates (Greenwald, Lettau and Ludvigson (2025)).

Relation to the Literature Our study builds on a large and growing literature modeling overreaction in subjective expectations and its relation to stock market behavior (Barberis, Shleifer and Vishny (1998), Chen, Da and Zhao (2013), Bordalo et al. (2018), Bordalo, Gennaioli, Ma and Shleifer (2020), Bordalo et al. (2019), Nagel and Xu (2022), Afrouzi, Kwon, Landier, Ma and Thesmar (2023), Bordalo et al. (2024), De La O and Meyers (2021, 2023) Hillenbrand and McCarthy (2021).) At the same time, other researchers argue that certain types of news are overlooked, causing underreaction (e.g., Mankiw and Reis (2002), Woodford (2002), Sims (2003), Gabaix (2019), Kohlhas and Walther (2021).) We extend these literatures by combining machine learning with structural estimation to freely estimate the direction and severity of a range of biases (if any) in the stock market’s reaction to hundreds of real-world news events, delineating the role of perceptions about multiple fundamental macro shocks in driving these reactions. Our findings add to those in the extant literature

by showing that markets can underreact to news even if investors overreact to all perceived shocks.

Other studies hypothesize that any link between subjectively expected future cash-flow growth and stock price variation occurs because the former responds to the latter rather than drives it (Bastianello and Fontanier (2022), Chaudhry (2023), Jin and Li (2023)) or, relatedly, that unexplained flows in and out of the stock market—disconnected from genuine cash-flow news—are responsible for stock market volatility (e.g., Gabaix and Koijen (2021), Hartzmark and Solomon (2022)). These papers study price movements driven by flows or mechanical factors unrelated to news, without taking a stand on what caused the price movement or flow to change in the first place. We take a converse yet complementary approach by studying market reactions to actual news, estimating its role in causing equilibrium price movements. Since actual news causes adjustments in forward-looking asset prices only when investors' subjective expectations are revised, such reactions are highly informative about investor beliefs.

We follow tradition by using equity analysts' survey forecasts of earnings growth as an observable indicator of subjective cash-flow expectations. As emphasized by Adam and Nagel (2023), however, the extent to which equity analysts' forecasts represent broader market expectations remains an open question. Our methodology takes a step toward addressing this limitation by employing a structural estimation that substantially broadens the set of observable indicators relevant for understanding investor beliefs. In our approach, the true underlying expectations of investors are identified by using a wide range of forward-looking indicators, including surveys and asset prices themselves, mapped onto theoretically motivated expressions that must obey cross-equation restrictions. This allows us to use multiple empirical signals to identify the subjective beliefs of stock market investors, going beyond the use of surveys alone.

The methodology of this paper builds off of the structural mixed-frequency approach of Bianchi, Ludvigson and Ma (2022b) (BLM2) for inferring what markets learn from news. Unlike the present study, BLM2 makes no use of machine learning to quantify systematic

expectational error; it investigates market reactions to news without addressing whether those reactions are nonrational and if so why. Merging machine learning and structural estimation fills this gap, providing an approach unique to the present paper and, to our knowledge, the extant literature.

The machine learning aspect of our methodology to measure systematic expectational errors consistent with the conditions of real-world expectation formation uses the general approach of BLM1 and Bianchi et al. (2025). The contribution of this paper is to take these machine-measured biases as an input into a structural estimation to investigate why those biases occur, with specific attention to how they show up in reactions to news. Our machine learning approach builds on insights in Bybee, Kelly, Manela and Xiu (2021), Gu, Kelly and Xiu (2020), and Cong, Tang, Wang and Zhang (2021), which show the power of supervised learning algorithms for asset return prediction. While our algorithms utilize supervised learning, they differ from these studies in that they are specifically designed to uncover and quantify distortions in subjective beliefs. A foundational principle of our algorithms recognizes that market participants have access to thousands of pieces of potentially relevant information in real time, while the canonical standard for efficient markets and rational expectation formation is predicated on the efficient use of all of it. The machine algorithm we design constructs a benchmark for objective expectation formation that is, by construction, free from human cognitive biases and efficiently copes with overfitting and structural change without look-ahead bias. Adherence to this principle is important to avoid overstating estimates of biases in the structural model.

The rest of this paper is organized as follows. In the next section we present a simplified framework to explain the key elements of our approach. We describe our machine learning algorithm in Section 3, the full structural model in Section 4, and the estimation, data, and measurement for the full structural model in Section 5. Section 6 presents our main findings. Section 7 presents additional results designed to unpack the main mechanisms behind our findings, while Section 8 concludes. Throughout the paper we use lowercase letters to denote log variables, i.e., $d_t = \ln(D_t)$, and “ $\sim$ ” to denote features of the model under the subjective

beliefs of the investor that may depart from full rationality.

2 Simplified Framework

This section contains two parts. The first part presents a simplified structural model of investor behavior and aggregate dynamics. The second part provides key steps of our empirical approach, which synthesizes the machine learning output with the structural estimation, using this simplified framework to illustrate the core elements of our approach. Since the application to the full structural framework is a straightforward generalization, we leave estimation details for the full structural model to the Internet Appendix.

Any marriage of machine learning with parametric structural estimation must confront the fact that the structural model is a stylized representation of reality subject to error, while the machine beliefs, survey forecasts, and other data are the product of much more complicated real-world phenomena. This section clarifies that our methodology produces results that are conditional on a stylized structural model, but one that we explicitly treat in estimation as an approximation of a complex and unknown “true” data generating process.

Simplified Structural Model Let real stock market payout, D_t , be a time-varying share K_t of real output Y_t , i.e., $D_t = K_t Y_t$. With arbitrary time-variation in K_t , the specification $D_t = K_t Y_t$ is a tautology. We argue, however, that log growth Δd_t is empirically better described by the specification $d_t = k_t + y_t$ than by a univariate process for Δd_t , because the former helps to identify distinct trend and cycle components that arise separately from variation in k_t and y_t . We present evidence on this below.

To see how these distinct components contribute short-run and longer-run components in earnings/payout growth, consider a simplified theoretical setting in which a representative investor forms subjective beliefs about log real stock market payouts, d , which follows the

law of motion:

$$\Delta d_t = \Delta y_t + k_t - k_{t-1} \quad (1)$$

$$k_t = (1 - \rho_k)k + \rho_k k_{t-1} + \varepsilon_{k,t} \quad (2)$$

$$\Delta y_t = (1 - \rho_{\Delta y})\Delta y + \rho_{\Delta y}\Delta y_{t-1} + \varepsilon_{\Delta y,t}. \quad (3)$$

Write the above as a bi-variate system in deviations from steady-state using “hats,” i.e.,

$\hat{k}_t \equiv k_t - k$:

$$\underbrace{\begin{bmatrix} \hat{\Delta d}_{t+1} \\ \hat{k}_{t+1} \\ \hat{\Delta y}_{t+1} \end{bmatrix}}_{\hat{S}_{t+1}^M} = \underbrace{\begin{bmatrix} 0 & \rho_k - 1 & \rho_{\Delta y} \\ 0 & \rho_k & 0 \\ 0 & 0 & \rho_{\Delta y} \end{bmatrix}}_{T^M(\theta^M)} \underbrace{\begin{bmatrix} \hat{\Delta d}_t \\ \hat{k}_t \\ \hat{\Delta y}_t \end{bmatrix}}_{\hat{S}_t^M} + \underbrace{\begin{bmatrix} 1 & 1 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}}_{R^M} \underbrace{\begin{bmatrix} \varepsilon_{k,t+1} \\ \varepsilon_{\Delta y,t+1} \end{bmatrix}}_{\varepsilon_{t+1}^M}, \quad (4)$$

or, letting $\theta^M \equiv (\rho_k, \rho_{\Delta y})'$, in matrix notation as

$$\hat{S}_{t+1}^M = T^M(\theta^M) \hat{S}_t^M + R^M \varepsilon_{t+1}^M. \quad (5)$$

Suppose both k_t and Δy_t are stationary with $0 \leq \rho_k, \rho_{\Delta y} < 1$. The system (4) therefore implies that expected log growth $\mathbb{E}_t[\Delta \hat{d}_{t+1}] = (\rho_k - 1)\hat{k}_t + \rho_{\Delta y}\hat{\Delta y}_t$ has both a negatively autocorrelated component originating from fluctuations in the payout share k_t , and a positively autocorrelated component originating from output growth Δy_t . It follows that a *negative* impulse to $\varepsilon_{k,t}$ generates the expectation of *positive* catch-up growth next period $\mathbb{E}_t[\Delta \hat{d}_{t+1}] > 0$, while a *positive* impulse to $\varepsilon_{k,t}$ generates the expectation of *negative* fall-back growth next period $\mathbb{E}_t[\Delta \hat{d}_{t+1}] < 0$. We refer to the k_t earnings share component as the “cyclical” component, to the Δy_t positively autocorrelated component as the “trend” component, and to the entire bi-variate specification as a “trend-cycle” model. This labeling serves to explicitly distinguish the bi-variate model from more commonly employed univari-

ate autoregressive models for Δd_t , such as those in Nagel and Xu (2022), and Bordalo et al. (2024).

Evidence for empirically relevant variation in the earnings share of output has previously been emphasized by Greenwald et al. (2025). Although earnings is not conceptually the same as payout, Greenwald et al. (2025) provide evidence consistent with the idea that it is the most important contributor to its variation over extended periods of time. For the full structural estimation, we use several noisy signals of aggregate payout, including aggregate earnings. In this section, we provide supporting evidence on the empirical relevance of earnings fluctuations around a trend, by reporting the results of specification tests comparing the fit of the trend-cycle specification (1)-(3) with that of a standard univariate autoregressive specification for $\Delta d_t = \mu + \rho \Delta d_{t-1} + \varepsilon_t$, using observations on earnings growth. For this purpose, we measure d_t with a bottom-up estimate of IBES “Street Earnings” for the S&P 500. Street Earnings differ from GAAP earnings by excluding discontinued operations, extraordinary charges, and other non-operating items. We discuss this measure of earnings further below. Table 1 provides the results of two specification comparisons, based on the estimated log likelihood and the BIC criterion, which varies inversely with the likelihood but penalizes for extra parameters.

Table 1: Model comparison on (street) earnings growth

	AR(1) with intercept	Trend-cycle
$\log L(\hat{\theta})$	-549.170	-544.759
BIC	1113.508	1104.687

ALT TEXT: Table comparing two earnings growth models based on log likelihood and the BIC.

By both measures, the data strongly prefer the trend-cycle model. This is key because, as shown below, modeling variation in k_t and y_t separately generates findings that differ markedly from univariate setups. Based on this evidence, we view the trend-cycle specification—and our subsequent findings—as superior to a standard univariate autoregressive specification for Δd_t .

We consider two types of distortion in investor beliefs about stock market fundamentals S_t^M . First, the perceived process for fundamentals growth may differ from (5) because the

investor's subjective value for θ^M differs from its objective value, resulting in:

$$\hat{S}_{t+1}^M = T^M \left(\tilde{\theta}^M \right) \hat{S}_t^M + R^M \tilde{\varepsilon}_{t+1}^M, \quad (6)$$

where $\tilde{\theta}^M \equiv (\tilde{\rho}_k, \tilde{\rho}_{\Delta y})'$ are the perceived parameters driving fundamentals persistence. Since the functional form of (6) is otherwise identical to that of (4), this implies distorted perceptions about θ^M translate directly into distorted perceptions about the shocks. In this case, the perceived shock vector $\tilde{\varepsilon}_t^M$ will differ from the objective innovation ε_t^M .²

Second, revisions in expectations may be subject to a time-varying distortion η_t . To model this distortion, we generalize the univariate specifications of Bordalo et al. (2018), Bordalo et al. (2019), and Bordalo et al. (2024) to accommodate the multivariate system (6). As in those specifications, investors are unaware that they have a distortion but behave *as if* their subjective expectations $\tilde{\mathbb{E}}_t[\cdot]$ were conditional on some additional news η_t . Whereas in those specifications η_t is a scalar, here it is a 3×1 vector satisfying

$$\tilde{\mathbb{E}}_t \left[\hat{S}_{t+1}^M \right] = T^M \left(\tilde{\theta}^M \right) \left(\hat{S}_t^M + \zeta \eta_t \right), \quad (7)$$

where $\eta_t \equiv (\eta_{\Delta d,t}, \eta_{k,t}, \eta_{\Delta y,t})'$. The scalar parameter ζ controls the magnitude and nature of the distortion and nests different models. If $\zeta > 0$, investor expectations overreact to their perceived news as in models with diagnostic expectations (DE) or earlier models of belief overreaction (e.g., Barberis et al. (1998)). If $\zeta < 0$, investors underreact to perceived news, as in models with inattention (Sims (2003), Gabaix (2019)). In the remainder of this paper, we refer to $\zeta \eta_t$ simply as the “DE distortion” for brevity, even though, strictly speaking, the reference to diagnostic expectations only applies when $\zeta > 0$. The empirical relevance of either type of distortion—captured by the sign and magnitude of ζ —will be subject to estimation in the full structural model.

As in Bordalo et al. (2018), Bordalo et al. (2019), and Bordalo et al. (2024), we allow

²“Peso problems” in which investors fear a rare event that does not occur in the sample observed by our machine algorithm would also show up as a wedge between $\tilde{\theta}^M$ and θ .

overreaction to gradually revert over time by specifying η_t to follow a VAR(1) (rather than AR(1)) process $\eta_t = \rho_\eta \tilde{T}^M \eta_{t-1} + R^M \tilde{\varepsilon}_t^M$, where $\tilde{T}^M \equiv T^M \left(\tilde{\theta}^M \right)$ and $0 \leq \rho_\eta < 1$. This shows that the distortion vector η_t has innovations that are proportional to the perceived cash-flow shocks. Note that when $\rho_\eta = 0$, $\eta_{\Delta d} = \tilde{\varepsilon}_{k,t} + \tilde{\varepsilon}_{\Delta y,t}$.

Taken together, the model specializes to rational expectations (RE), or objective beliefs, when $\theta^M = \tilde{\theta}^M$ and $\zeta = 0$.

Two points about this specification are worthy of emphasis. First, while ζ is a scalar parameter that is the same for all shocks, equation (7) shows that the overall DE distortion is the *product* $\zeta \eta_t$. This demonstrates that there is a shock-specific component to the dynamic DE distortion. Shocks that are more volatile or more persistent will generate larger distortions even if the same ζ applies to all shocks. Second, (6) implies that the investor may misperceive the law of motion for cash-flow growth. This does not mean that investors suffer from delusions about realized cash-flow growth Δd_{t+1} , once observed. What the distinction between (5) and (6) does imply is that investors may disagree with a fully rational agent about how they got to that realization.

Suppose for now that investors price in a constant risk-premium and risk-free rate r_f under their subjective beliefs. (The full model relaxes this assumption.) Let P_t^D denote the stock price level and apply a Campbell and Shiller (1989) approximate present value identity by expanding the log return $r_{t+1}^D \equiv \ln(P_{t+1}^D + D_{t+1}) - \ln(P_t^D)$ around a point $P_t^D/D_t \equiv PD$:

$$r_{t+1}^D = \kappa_{pd,0} + \beta p d_{t+1} - p d_t + \Delta d_{t+1}, \quad (8)$$

where r_t^D is the stock market return, $p d_t \equiv p_t^D - d_t$, $\beta \equiv \frac{PD}{1+PD}$, $\kappa_{pd,0} \equiv \ln(1 + \exp(pd)) - \beta(pd)$, and $pd = \ln(PD)$. With the constant risk-premium and risk-free rate and imposing

$\lim_{j \rightarrow \infty} \beta^j pd_{t+j} = 0$, the price-payout ratio is

$$pd_t = pd + \tilde{\mathbb{E}}_t \sum_{v=0}^{\infty} \beta^v \hat{\Delta} d_{t+1+v} \quad (9)$$

$$= pd + \left(\frac{\tilde{\rho}_{\Delta y}}{1 - \tilde{\rho}_{\Delta y} \beta} \right) \left(\hat{\Delta} y_t + \zeta \eta_{\Delta y, t} \right) + \left(\frac{\tilde{\rho}_k - 1}{1 - \tilde{\rho}_k \beta} \right) \left(\hat{k}_t + \zeta \eta_{k, t} \right), \quad (10)$$

where $pd \equiv (\kappa_{pd,0} - r^D + \Delta d) / (1 - \beta)$, r^D equals to the constant subjectively expected return, and $\Delta d = \Delta y$ equals steady-state payout growth. This shows that the price-payout ratio reflects the investor's perceived law of motion and the time-varying distortion η_t . Combining (10) and (8) and noting that $\tilde{\mathbb{E}}_t [\eta_{t+1}] \equiv 0$ because the agent is unaware of the distortion, we verify $\tilde{\mathbb{E}}_t [r_{t+1}^D] = r^D$.

If this were the model that generated the data, what would be the objective (i.e., non-distorted) expectation of future returns? In contrast to (7), objective beliefs take the form

$$\mathbb{E}_t [\hat{S}_{t+1}^M] = T^M (\theta^M) \hat{S}_t^M.$$

Taking expectations of (8) under objective beliefs $\mathbb{E}_t[\cdot]$ yields

$$\begin{aligned} \mathbb{E}_t [r_{t+1}^D] = r^D &+ \left[\frac{\rho_{\Delta y} - \tilde{\rho}_{\Delta y}}{1 - \tilde{\rho}_{\Delta y} \beta} \right] \hat{\Delta} y_t + \left[\frac{\beta \rho_{\Delta y} \rho_{\eta} - 1}{1 - \tilde{\rho}_{\Delta y} \beta} \right] \tilde{\rho}_{\Delta y} \zeta \eta_{\Delta y, t} \\ &+ \left[\frac{(\rho_k - \tilde{\rho}_k)(1 - \beta)}{1 - \tilde{\rho}_k \beta} \right] \hat{k}_t + \left[\frac{\beta \rho_k \rho_{\eta} - 1}{1 - \tilde{\rho}_k \beta} \right] (\tilde{\rho}_k - 1) \zeta \eta_{k, t} \end{aligned} \quad (11)$$

Subjective and objective expected returns coincide only when (i) $\zeta = 0$, and (ii) $\tilde{T}^M = T^M$, in which case objective expected returns in (11) are always r^D , and investors rationally price in a constant risk-free rate and risk premium. More generally, the terms in square brackets show the predictable components of objectively expected future returns that are attributable to the systematic distortions in subjective beliefs.

Belief Reactions to News: The Shock Composition Effect With these expressions in hand, we now consider how different news events would affect subjective and objective

beliefs, where a news “event” in this context is defined as a distinct combination of perceived economic shocks, or revisions to subjective expectations about the current economic state. For ease of exposition, we set $\rho_\eta = 0$ in (7), implying $\eta_{k,t} = \tilde{\varepsilon}_{k,t}$, $\eta_{\Delta y,t} = \tilde{\varepsilon}_{\Delta y,t}$.

1. Event 1: $\tilde{\varepsilon}_{k,t} < 0$. This news causes the investor to revise her perception of the current payout share downward, while having no effect on perceived output growth, i.e., $\tilde{\varepsilon}_{\Delta y,t} = 0$.

- (a) Suppose $\zeta > 0$ as in DE models of belief overreaction and let $\tilde{\rho}_k = \rho_k = 0$. It is straightforward to show that investors respond to this news with excessive optimism about catch-up growth in payout:

$$\left(\tilde{\mathbb{E}}_t - \mathbb{E}_t\right) \left[\hat{\Delta}d_{t+1}\right] = -\zeta\tilde{\varepsilon}_{k,t} > 0,$$

since $\tilde{\varepsilon}_{k,t} < 0$. The excessive optimism inflates the initial price impact, but from (11) we can see that the inevitable investor disappointment in future growth (once observed) will cause a price reversal and lower future returns that is objectively predictable: $\mathbb{E}_t[r_{t+1}^D] = r^D + \zeta\tilde{\varepsilon}_{k,t} < r^D$.

- (b) Suppose $\zeta = 0$ while $\tilde{\rho}_k > \rho_k$. Here the investor over-extrapolates today’s bad news to the future, generating excessive pessimism about future growth:

$$\left(\tilde{\mathbb{E}}_t - \mathbb{E}_t\right) \left[\hat{\Delta}d_{t+1}\right] = (\tilde{\rho}_k - \rho_k)\tilde{\varepsilon}_{k,t} < 0.$$

The excessive pessimism means that the investor will inevitably be favorably surprised in the future, causing a price rebound and objective expectation of higher future returns: $\mathbb{E}_t[r_{t+1}^D] = r^D + [(\rho_k - \tilde{\rho}_k)(1 - \beta)/(1 - \tilde{\rho}_k\beta)]\tilde{\varepsilon}_{k,t} > r^D$.

2. Event 2: $\tilde{\varepsilon}_{\Delta y,t} < 0$. This news causes the investor to revise her perception of current output growth downward, while having no effect on the perceived k_t , i.e., $\tilde{\varepsilon}_{k,t} = 0$.

- (a) Suppose $\zeta > 0$ and $\tilde{\rho}_{\Delta y} = \rho_{\Delta y} > 0$.³ Investors respond with excessive pessimism about subsequent growth: $(\tilde{\mathbb{E}}_t - \mathbb{E}_t) [\hat{\Delta}d_{t+1}] = \tilde{\rho}_{\Delta y} \zeta \tilde{\varepsilon}_{\Delta y, t} < 0$, since $\tilde{\varepsilon}_{\Delta y, t} < 0$. This causes the price to overreact on the downside, which (11) shows leads to a predictable price reversal and objective expectation of higher future returns.
- (b) Suppose $\zeta = 0$, while $\tilde{\rho}_{\Delta y} > \rho_{\Delta y}$. The investor over-extrapolates today's bad economic growth news to the future, generating excessive pessimism about subsequent growth: $(\tilde{\mathbb{E}}_t - \mathbb{E}_t) [\hat{\Delta}d_{t+1}] = (\tilde{\rho}_{\Delta y} - \rho_{\Delta y}) \tilde{\varepsilon}_{\Delta y, t} < 0$. Like 2 (a), the price overreacts on the downside, generating a predictable price reversal and objective expectation of higher future returns.
3. Event 3: $\tilde{\varepsilon}_{k, t} < 0$ and $\tilde{\varepsilon}_{\Delta y, t} < 0$. This news causes investors to revise their subjective expectation of both the payout share and output growth downward. Computing the overall market impact of this news requires combining the reactions to both perceived shocks. For a concrete numerical example, we consider the situation where $\zeta > 0$ and the other conditions of Cases 1(a) and 2(a) apply. From 1(a), DE causes investors to respond to $\tilde{\varepsilon}_{k, t} < 0$ with excessive optimism about catch-up growth. Consequently, under these distorted beliefs, p_t^D would rise by some amount (suppose 5), whereas under RE with $\rho_k = 0$ the ex-dividend price would be unchanged. From 2(a), DE causes investors to respond to $\tilde{\varepsilon}_{\Delta y, t} < 0$ with excessive pessimism about future growth. Meanwhile, under these same distorted beliefs, prices would fall by some amount (suppose 10), whereas under RE with $\rho_{\Delta y} > 0$ prices would fall by a lesser degree (suppose 6). Taken together, the behavioral model implies an overall price impact of $5 - 10 = -5$, which can be compared to the impact under objective beliefs of $0 - 6 = -6$. Thus, the market *underreacts* to the news as a whole, even though the investor *overreacts* to all shocks.

Event 3 gives rise to an important distinction between the multivariate setting studied here and the DE models typical of the literature, in which $\zeta > 0$ applies to a univariate earnings or payout process. When multiple primitive macroeconomic risks are relevant for the subjective

³This example requires $\tilde{\rho}_{\Delta y} \neq 0$ because DE operates only on innovations that have predictability for future growth.

growth expectations that underpin shareholder value, overreaction to all shocks can dampen rather than amplify market volatility via a *shock composition effect*. This happens when news events cause investors to revise their perceptions about more than one fundamental shock, in directions that have counteracting but *asymmetric* implications for valuations. In the example of Event 3 above, the market fell “too little” because the investor’s revised expectations for the earnings share were more overly rosy than her views on economic growth were overly pessimistic. Although this example plugs in hypothetical values for the price effects, it serves to illustrate the point that asymmetries can arise even though the same $\zeta > 0$ scalar parameter applies to all shocks because, as (7) shows, the shock-specific volatility and propagation properties still matter for the magnitude of shock-specific overreactions.⁴ For the main application of this paper, the volatility and propagation parameters governing these asymmetries are estimated and the extent to which such asymmetric overreactions play a role in historical stock market variation is a key empirical question to be explored.

Estimation We estimate the model using Bayesian state-space methods. A general premise of the approach is that a wide variety of observable data—interpreted through the lens of a structural model—constitute important signals of what real-world market participants believe and expect. These include direct measures of subjective asset market expectations from surveys of equity analysts and investors (as in the traditional approach), as well as fluctuations in spot prices, futures markets, and professional forecasts of the broader economy. Below we explain how the structural estimation uses the machine expectation output, denoted $\mathbb{E}_t^{ML}[\cdot]$, as noisy signals of the theoretical RE expectation $\mathbb{E}_t[\cdot]$, a procedure that forces our estimates of $\mathbb{E}_t[\cdot]$ to be consistent with a real-time objective expectation process.

To illustrate the procedure, we begin with the solution to the model (5)-(10), which implies that the state vector $S_t = [S_t^M, pd_t, pd_{t-1}, r_t^D, \eta_t, S_t^*]'$ evolves according to a vector

⁴This can be observed from (7) by noting that η_t contains the perceived shocks and multiplies $T^M(\tilde{\theta}^M)$, which contains the propagation parameters.

autoregression (VAR) state equation

$$S_t = C(\Theta) + T(\Theta)S_{t-1} + R(\Theta)Q\varepsilon_t^M,$$

where $S_t^M = (\Delta d_t, k_t, \Delta y_t)$ is a set of macro fundamentals, S_t^* is discussed below, C , T , and R are matrices composed of the model's primitive parameters

$$\Theta = (\rho_k, \rho_{\Delta y}, \tilde{\rho}_k, \tilde{\rho}_{\Delta y}, \zeta, r^D, \beta, \rho_\eta, k, \Delta y)',$$

Q is a matrix of shock volatilities, and η_t is the latent DE distortion to be estimated. The relation between the variables in the model and a vector of observable signals X_t can be written as an observation equation taking the form:

$$X_t = D + ZS_t + Uv_t,$$

where D and Z are matrix parameters, and v_t is a vector of observation errors with standard deviations in the diagonal matrix U . The observation errors v_t are important for modeling noise due to various sources (including gaps that arise from the approximating structural model itself) and are discussed below. Combining $X_t = D + ZS_t + Uv_t$ with the state equation $S_t = C + TS_{t-1} + RQ\varepsilon_t^M$, allows us to estimate the model parameters and theoretical states S_t using state-space methods.

In the model description above, investor expectations were conditioned on the state vector S_t . In reality, however, some of its elements will be observed imperfectly in real time because they undergo subsequent revision. For example, asset price data p_t^D are not subject to revision, but d_t is *real* payout and must be computed using data on inflation that is subject to revision. To better match the conditions of real-world decision making, in our estimation we assume that investors have access only to a noisy measure of any indicators subject to subsequent revision, and price assets on that basis. Let $S_t^* = [\Delta d_t^*, k_t^*, \Delta y_t^*, pd_t^*, pd_{t-1}^*, r_t^{D*}]$ denote these noisy elements of S_t observed in real time.

Let $\mathbb{F}_t[y_{t+v}]$ generically denote a vector of observed subjective forecast measures made at time t of variable y at time $t+v$ measured from surveys, futures markets, or other expectations data, and let $\mathbb{E}_t^{ML}[y_{t+v}]$ denote an observed objective machine forecast produced in an outer estimation. Let matrices with a subscript, e.g., Z_x , denote the parameter sub-vector of Z that when multiplied by S_t or S_t^* and added to $D_x + U_x v_{x,t}$ picks out the appropriate model variable to map back into empirical observations X_t , e.g., $Z_k S_t$ picks out the element of S_t corresponding to k_t . Finally, collect the coefficients on $\hat{k}_t$, $\hat{y}_t$, and $(\eta_{\Delta y,t}, \eta_{k,t})'$ in (11) showing the effect of these variables on objective return expectations $\mathbb{E}_t[r_{t+1}^D]$ into $Z_{\mathbb{E}(r),k}$, $Z_{\mathbb{E}(r),\Delta y}$, and $Z_{\mathbb{E}(r),\eta}$, respectively. The observation equation $X_t = D + ZS_t + Uv_t$ takes the form

$$\begin{bmatrix} [\Delta d_t, k_t, \Delta y_t]' \\ pd_t \\ r_t^D \\ \mathbb{F}_t[\hat{\Delta d}_{t+1}] \\ \mathbb{E}_t^{ML}[\hat{\Delta d}_{t+1}] \\ \mathbb{F}_t[r_{t+1}^D] \\ \mathbb{E}_t^{ML}[r_{t+1}^D] \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ r^D \\ r^D \end{bmatrix} + \begin{bmatrix} Z_{sM} S_t \\ Z_{pd} S_t \\ Z_r S_t \\ ((\tilde{\rho}_k - 1)Z_k + \tilde{\rho}_{\Delta y} Z_{\Delta y} + \zeta Z_\eta) S_t^* \\ ((\rho_k - 1)Z_k + \rho_{\Delta y} Z_{\Delta y}) S_t^* \\ 0 \\ (Z_{\mathbb{E}(r),k} + Z_{\mathbb{E}(r),\Delta y} + \zeta Z_{\mathbb{E}(r),\eta}) S_t^* \end{bmatrix} + Uv_t. \quad (12)$$

The vector on the left consists of empirical observations. The vector with theoretical states on the right must obey the cross-equation restrictions implied by (5)-(10). The above mapping between observations and model restrictions illustrates key steps of our estimation approach.

1. **Historical data** $[\Delta d_t, k_t, \Delta y_t]'$ are mapped onto the model's approximating objective laws of motion $Z_{sM} S_t$ to obtain best-fitting descriptions of structural model data dynamics. Observation errors in v_t account for both estimation and specification error arising because the model is an approximation of the true (unknown) data dynamics.
2. **Real time data.** Investors and machine forecasts are made on the basis of *real time* data S_t^* , as in real-world forecasting.

3. **Multiple signals identify $\tilde{\mathbb{E}}_t[\cdot]$.** Multiple forward-looking indicators are used to identify subjective beliefs, including financial market variables (e.g., r_t^D and pd_t), survey and futures markets forecasts $\mathbb{F}_t[\cdot]$, each of which is treated as noisy signals of the true underlying subjective expectations process $\tilde{\mathbb{E}}_t[\cdot]$ of the investor.⁵ For example, multiple subjective expectations measures $\mathbb{F}_t[\hat{\Delta}d_{t+1}]$ map into $((\tilde{\rho}_k - 1)Z_k + \tilde{\rho}_{\Delta y}Z_{\Delta y} + \zeta Z_\eta) S_t^*$, informing estimates of $\tilde{\rho}_k$, $\tilde{\rho}_{\Delta y}$, and ζ . As we often have multiple noisy signals of a single theoretical concept, observation error is inevitable.
4. **Machine forecasts identify $\mathbb{E}_t[\cdot]$.** Iterative machine forecasts $\mathbb{E}_t^{ML}[\cdot]$ serve as noisy signals of the theoretical RE benchmark, e.g., $\mathbb{E}_t^{ML}[\hat{\Delta}d_{t+1}]$ is mapped onto $\mathbb{E}_t[\hat{\Delta}d_{t+1}] = ((\rho_k - 1)Z_k + \rho_{\Delta y}Z_{\Delta y}S_t^*)$. This forces our structural estimates of $\mathbb{E}_t[\cdot]$ to be consistent with a real-time objective expectations process based on knowledge we can verify would have been available to investors in real-time.

At the core of this approach is a strategy for using information from high dimensional, nonparametric, machine-based representations of objective beliefs to inform and identify systematic expectational errors as represented in stylized parametric frameworks of human behavior. There are three components to the approach: (i) the machine forecasts $\mathbb{E}_t^{ML}[\cdot]$ that are used as an empirical signal of objective beliefs, (ii) the systematic expectational errors (if any) that are embedded in observed investor behavior, and (iii) the stylized parametric framework of human behavior that the observations map onto.

For the first component, two aspects are central to our approach. First, the real-time nature of the machine estimation is designed to emulate the real-world setting and eliminate look-ahead advantages. Second, a high-dimensional neural network function serves to approximate what is ultimately the unknown function that best represents objective beliefs.⁶

⁵Short samples for survey expectations or other data are not technically a problem for this methodology since a much larger set of observables is used to measure expectations while missing values can be estimated using a filter and structural model.

⁶It is known that a multi-layer neural network can approximate virtually any unknown function arbitrarily well given a large enough set of inputs. See Hecht-Nielsen (1987) for the well-known Kolmogorov universal representation theorem that applies to arbitrary continuous functions and Ismailov (2023) for the theorem extending to discontinuous functions.

These empirically optimal forecasts are then mapped onto the appropriate equations of a low-dimensional structural model, providing approximately unbiased signals of what could have been rationally expected in the model environment.

The second component refers to the multiple real-world signals of investor expectations, including from asset prices themselves, that contribute to evidence of distortion. However, both machine and investors must cope with estimation, specification, and data-revision errors, as well as with structural change in an evolving environment. These aspects represent noise that are common to rational and subjective beliefs and that we accommodate in the structural estimation by allowing for errors in the observation equations of the state-space representation.

The third component refers to the primary purpose of structural modeling, which is to provide a conceptual framework for interpreting the data. Such frameworks are always approximations of reality, but help us relate findings to an existing literature, while making theoretical concepts precise and facilitating understanding.

Putting this all together, conditional on a stylized parametric model, the estimation procedure uses surveys and other forward-looking data to inform subjective parameters and distortions, machine forecasts inform objective beliefs, and observation errors capture noise.⁷

It is important to clarify that, in undertaking this approach, we do not assume that the machine beliefs $\mathbb{E}_t^{ML}[\cdot]$ —generated using financial market data—are the same as the beliefs of an agent in a counterfactual rational expectations model. Rather, it adopts the RE requirement—efficiently processing all available information—from the perspective of an unbiased forecaster observing an economy where real-world agents may be nonrational. Under this formulation, a machine algorithm (without human bias) that takes as inputs both objective economic data and subjective market data generated by agents with bias can learn

⁷In the first equations of (12) the structural model laws of motion are mapped back into full revised, historical data. These equations could be dropped from the estimation, so that the machine forecasts are the only signal on the parameters of the objective laws of motion. The cost of doing so is that these mappings are likely to improve the description of the model's historical relationships. Keeping them allows the estimator to strike a balance between doing a good job of describing such dynamics (as in traditional structural estimation), while at the same time mitigating concerns about overfitting and look-ahead bias that can arise from a purely in-sample structural estimation.

to exploit the market's systematic mistakes to better predict returns even if they wouldn't be predictable were agents fully rational. This phenomenon is exactly what is illustrated by equation (11). Consistent with this, the machine forecasts $\mathbb{E}_t^{ML}[\cdot]$ of fundamentals are mapped onto the objective laws of motion of macro fundamentals S^M , which hold under the parameter restrictions $\tilde{T}^M = T^M$, while the machine forecasts $\mathbb{E}_t^{ML}[\cdot]$ of financial market returns are mapped onto the laws of motion for subjective expected returns, which take into account that investors may hold beliefs characterized by $\zeta \neq 0$, and $\tilde{T}^M \neq T^M$. Once the parameters ζ and $\tilde{T}^M$ are identified via structural estimation, we can conduct true RE counterfactuals in the model, in which $\zeta \equiv 0$, and $\tilde{T}^M \equiv T^M$. These true counterfactuals are not the same as the empirical objects $\mathbb{E}_t^{ML}[\cdot]$ required to identify how objective beliefs differ from subjective beliefs.

The final step in the empirical analysis is to measure market reactions to news, which we do by employing the mixed-frequency filtering algorithm developed in BLM2 to estimate revisions in investor perceptions in tight windows surrounding news events. The nature and severity of any behavioral biases in market reactions to news is estimated by comparing jumps in model-implied investor beliefs with those of a counterfactual investor with rational expectations, whose beliefs are informed by our machine learning output. This leads us to discuss machine beliefs, which are compiled from algorithmic output and produced in a first-stage for use in X_t .

3 Machine Learning

To measure distortions in beliefs, we require a practical measure of unbiased, information-efficient expectation formation under the conditions of real-world decision making, with which to compare the subjective beliefs of investors. For this, we make use of the machine learning algorithms BLM1 and Bianchi et al. (2025). The contribution of these papers is to measure the overall magnitude of these distortions. The contribution of this paper is to take these machine-measured distortions as an input into a structural estimation in order

to investigate *why* those biases occur, with specific attention paid to how they show up in reactions to news. We refer the reader to BLM1 and Bianchi et al. (2025) (BLLM) for additional details on the machine estimation and output, providing only a summary description here.

We are interested in forming a machine expectation of a time series $y_{j,t+v}$ indexed by j whose value in period $v \geq 1$ the machine is asked to predict. The following machine specification is estimated over rolling samples:

$$y_{j,t+v} = G^e(\mathcal{X}_t, \beta_{j,v,t}) + \epsilon_{jt+v}. \quad (13)$$

where $\mathcal{X}_t$ is a large input dataset available in real time, including an intercept, and $G^e(\cdot)$ is a machine learning estimator that can be represented by a high dimensional set of finite-valued parameters $\beta_{j,v,t}$.⁸ With this estimator in hand, we follow the six step algorithmic approach of BLM1: 1. Sample partitioning,⁹ 2. Training, 3. Model selection and cross-validation, 4. Grid and sample partition re-optimization, 5. Out-of-sample prediction, 6. Roll forward and repeat. Step 3 includes variable selection, shrinkage, and hyper-parameter tuning. The end product of this procedure is a time-series of objective time t machine “beliefs” about $y_{j,t+v}$, denoted $\mathbb{E}_t^{ML}[y_{j,t+v}]$.

Two points about the algorithm bear emphasis. First, the machine expectations are based on only that information at t that we can verify would have been available to investors in real time. Second, the machine algorithm is designed to uncover *bias* in subjective beliefs, i.e., *predictable* mistakes that arise from a demonstrable misuse of available information. In

⁸We use the Long Short-Term Memory (LSTM) deep sequence recurrent neural network estimator with N hidden layers $h_t^n \in \mathbb{R}^{D_{hn}}$

$$G^{LSTM}(\mathcal{X}_t, \theta_{jh}) = \sum_{n=1}^N \underbrace{W^{(yh^n)}}_{1 \times D_{hn}} \underbrace{h_t^n(\mathcal{X}_t, \theta_{jh})}_{D_{hn} \times 1} + \underbrace{b_y}_{1 \times 1}.$$

⁹At time t , a prior training sample of size $\dot{T}$ is partitioned into two subsample windows: an “estimation” subsample consisting of the first T_E observations, and a hold-out “validation” sample of T_V subsequent observations so that $\dot{T} = T_E + T_V$.

the estimation below, surveys are used as signals of subjective beliefs. The algorithms of BLM1 and BLLM are structured so that the machine's forecasts can differ from the survey forecasts only if the machine finds evidence of predictable mistakes in the survey responses immediately prior to the machine making a true out-of-sample forecast. These algorithms are run multiple times while being "paired" with a different survey forecast, to identify predictable mistakes in every survey response.

The output of BLM1 and BLLM shows that the machine achieves sizable reductions in the mean-square-forecast-errors relative to survey forecasts over an extended testing subsample for stock market returns, earnings growth, output growth, and inflation. These reductions are largest during times of important economic change (see the papers for details). Overall, these results are consistent with the premise that a relatively unbiased, information-efficient machine using only real-time information is able to detect patterns in widely available data that notably improve predictive accuracy over human forecasts. This systematically superior performance motivates our use of the machine benchmark for measuring non-distorted expectation formation in the structural estimation. It is noteworthy, as shown in Bianchi et al. (2025), that survey respondents make much larger systematically predictable errors in their forecasts of earnings growth than in their forecasts of broad economic growth or returns. Keeping in mind that the machine forecasts are a central data input into the structural estimation identifying objective expectations, these preliminary results foreshadow and help explain a key finding below, namely that investor reactions to payout-share shocks are more distorted than they are to other shocks.

One might reasonably ask why this machine component is needed at all. After all, an alternative would be to construct a RE benchmark by estimating a presumed structural model of objective beliefs on historical data. The difficulty with this approach is three-fold. First, it is both subject to look-ahead bias and presumes perfect knowledge of the data generating process, factors that tend to overstate behavioral biases (BLM1, Farmer, Nakamura and Steinsson (2024)). Second, such an approach is silent on the cumulative importance of distortions beyond those implied by the chosen parametric model. Addressing this gap

requires an explicit measure of non-distorted expectation formation, against which we can measure behavioral distortions in the structural model. Third, parametric models may not be flexible enough to approximate the decision making of financial market participants. A machine algorithm can be highly flexible, while selecting the optimal amount of sparsity and shrinkage.

4 Structural Model

We now apply the ideas presented above for the simplified model to the full structural model. We work with a risk-adjusted log-linear approximation to the model, in which all random variables are conditionally log-normally distributed.

Macro Dynamics As above, let aggregate stock market payout, D_t , be a time-varying share K_t of real output Y_t , i.e., $D_t = K_t Y_t$. We now generalize the simple bivariate process considered above to allow for additional variables, each with their own short- and longer-run components. Specifically, macro dynamics are described by a series of equations for the nominal short rate i_t , general price inflation π_t , output growth Δy_t , and the log payout share of output $k_t \equiv d_t - y_t$. For each of these, we specify “trend” or “long-run” components denoted with “bars” that evolve according to

$$\bar{x}_t = (1 - \phi_x)\bar{x}_{t-1} + \phi_x x_t + \sigma_{\bar{x}, \xi_t} \varepsilon_{\bar{x}, t}, \quad \forall x = \{i, \pi, \Delta y, k\}, \quad (14)$$

where $\varepsilon_{\bar{x}, t} \sim \mathcal{N}(0, 1)$ is an i.i.d. shock to the trend component of x with a time-varying volatility $\sigma_{\bar{x}, \xi_t}$ discussed below, and ϕ_x is a parameter governing its persistence. We assume that i_t , π_t , Δy_t , and k_t vary cyclically around these trend components in the following manner:

The nominal short rate is presumed to be set by the central bank and follows the process

$$i_t - i = (1 - \psi_i) [\psi_\pi (\bar{\pi}_t - \pi) + \psi_{\Delta y} (\bar{\Delta y}_t - g)] + \psi_i (\bar{i}_{t-1} - i) + \sigma_{i, \xi_t} \varepsilon_{i, t}, \quad (15)$$

where $\varepsilon_{i,t} \sim \mathcal{N}(0, 1)$ is an i.i.d. monetary policy shock, and i, π , and g are parameters. The dynamics of inflation and output growth follow similar primitive processes:

$$\pi_t - \pi = \beta_{\pi,\pi} (\bar{\pi}_{t-1} - \pi) + \beta_{\pi,\Delta y} (\bar{\Delta y}_t - g) + \beta_{\pi,i} (\bar{i}_{t-1} - i) + \sigma_{\pi,\xi_t} \varepsilon_{\pi,t} \quad (16)$$

$$\Delta y_t - g = \beta_{\Delta y,\pi} (\bar{\pi}_{t-1} - \pi) + \beta_{\Delta y,\Delta y} (\bar{\Delta y}_{t-1} - g) + \beta_{\Delta y,i} (\bar{i}_{t-1} - i) + \sigma_{\Delta y,\xi_t} \varepsilon_{\Delta y,t}, \quad (17)$$

where $\beta_{i,j}$ are parameters and $\varepsilon_{\pi,t} \sim \mathcal{N}(0, 1)$ and $\varepsilon_{\Delta y,t} \sim \mathcal{N}(0, 1)$ are i.i.d. shocks that represents short-run, cyclical variation in these variables. The log payout share, k_t , is modeled as a primitive process as follows:

$$k_t - k = \rho_{k,k} (\bar{k}_{t-1} - k) + \beta_{k,\Delta y} (\bar{\Delta y}_t - g) + \sigma_{k,\xi_t} \varepsilon_{k,t}, \quad (18)$$

where $\varepsilon_{k,t} \sim \mathcal{N}(0, 1)$ is an i.i.d. shock. This specification implies that inflation, output growth, the payout share, and the short-rate are simultaneously determined by the dynamical system (15)-(18).¹⁰

We refer to i.i.d. innovations without the bars as the *cyclical* components (e.g., $\varepsilon_{k,t}$ is the “cyclical payout share shock”), to those in (14) as *trend* component shocks (e.g., $\varepsilon_{\bar{k},t}$ is the “trend payout share shock”), and to the overall specification as a *trend-cycle model*, a generalization of the simplified trend-cycle specification. It should be kept in mind, however, that the “trend” components are latent random variables that are hybrids of i.i.d. and persistent processes and are furthermore contemporaneously correlated with multiple economic variables in the simultaneous system above. We use these hybrid specifications to introduce parsimoniously parameterized but flexible persistence in the variables in a manner similar to a vector autoregression, but with fewer estimable parameters.

The shock volatilities in all of primitive processes above vary with the discrete valued random variable ξ_t , which evolves according to a $\mathcal{N}$ -state Markov-switching process with transition matrix $\mathbf{H}$. Collect the parameters $\psi_i, \phi_\pi, \dots$ etc., of the above equations including

¹⁰This specification for macro dynamics is consistent with a triangular identification strategy for monetary policy shocks.

$\mathbf{H}$ into a vector θ^M . Equations (15)-(18), along with the expression for payout growth, $\Delta d_t = \Delta k_t + \Delta y_t$, represent a macro-dynamic system that can be expressed as a Markov-switching vector autoregression (MS-VAR) law of motion (LOM) taking the form:

$$S_t^M = C^M(\theta^M) + T^M(\theta^M)S_{t-1}^M + R^M(\theta^M)Q_{\xi_t}^M\varepsilon_t^M, \quad (19)$$

where $S_t^M \equiv [\Delta y_t, \overline{\Delta y}_t, \Delta d_t, \pi_t, \overline{\pi}_t, i_t, \overline{i}_t, k_t, \overline{k}_t]'$, $C^M(\cdot)$, $T^M(\cdot)$, $R^M(\cdot)$ are matrices of primitive parameters θ^M , $\varepsilon_t^M = [\varepsilon_{\Delta y,t}, \varepsilon_{\overline{\Delta y},t}, \varepsilon_{\pi,t}, \varepsilon_{\overline{\pi},t}, \varepsilon_{i,t}, \varepsilon_{\overline{i},t}, \varepsilon_{k,t}, \varepsilon_{\overline{k},t}]'$ is a vector of primitive macro shocks, and $Q_{\xi_t}^M(\cdot)$ is a diagonal matrix of shock volatilities that varies stochastically with ξ_t . Due to the endogeneity of these variables, $R^M(\cdot)$ has non-zero off diagonal elements, implying that multiple fundamental shocks affect a single state variable.

Perceived Macro Dynamics Investors have subjective beliefs $\tilde{\theta}^M$ about the parameters governing macro dynamics in (15)-(18) that could differ from the objective θ^M . Let these differences be captured by a wedge vector w_θ : $\tilde{\theta}^M = \theta^M + w_\theta$. We assume that investors apply these perceived dynamics to a noisy measure of S_t^M that they observe in real time, denoted S_t^{M*} . The two are related by $A_o S_t^{M*} = A_o S_t^M + Q^v \varepsilon_{v,t}$, where $\varepsilon_{v,t} \sim \mathcal{N}(0, 1)$ is an i.i.d. “vintage” error attributable to data revisions.¹¹ Elements of S_t^M for which there is no post-publication revision are assumed to have no such vintage errors. Investors take S_t^{M*} as given and price assets accordingly.¹² Taken together, these assumptions imply that the perceived counterpart to (19) takes the form

$$S_t^{M*} = C^M(\tilde{\theta}^M) + T^M(\tilde{\theta}^M)S_{t-1}^{M*} + R^M(\tilde{\theta}^M)\tilde{Q}_{\xi_t}^M\tilde{\varepsilon}_t^M \quad (20)$$

$$S_t^{M*} \equiv [\Delta y_t^*, \overline{\Delta y}_t^*, \Delta d_t^*, \pi_t^*, \overline{\pi}_t^*, i_t^*, \overline{i}_t^*, k_t^*, \overline{k}_t^*]' \quad (21)$$

$$\tilde{\varepsilon}_t^M \equiv [\tilde{\varepsilon}_{\Delta y,t}, \tilde{\varepsilon}_{\overline{\Delta y},t}, \tilde{\varepsilon}_{\pi,t}, \tilde{\varepsilon}_{\overline{\pi},t}, \tilde{\varepsilon}_{i,t}, \tilde{\varepsilon}_{\overline{i},t}, \tilde{\varepsilon}_{k,t}, \tilde{\varepsilon}_{\overline{k},t}]', \quad (22)$$

¹¹The A_o matrix emphasizes that vintage errors can be on a linear combination of elements of S_t^{M*} and/or that they apply only to specific elements.

¹²This treats S_t^{M*} as an unbiased signal of the underlying “true” state vector S_t^M that is precise enough to reasonably ignore any uncertainty about the signal when pricing assets.

where $\tilde{\varepsilon}_t^M$ is a vector of perceived primitive macroeconomic shocks. The perceived volatilities $\tilde{Q}_{\xi_t}^M$ of these shocks vary with the same discrete valued random variable ξ_t but have a perceived transition matrix $\tilde{\mathbf{H}}$ that may differ from $\mathbf{H}$. As in the simplified model, $\tilde{R}^M$ is neither square nor diagonal, so distorted beliefs about the parameters translate directly into distorted perceptions about the shocks, implying that investors can misattribute a change in one primitive shock to a mixture of others.

Let $\tilde{T}^M \equiv T^M(\tilde{\theta}^M)$ and analogously for $\tilde{R}^M(\tilde{\theta}^M)$ and $\tilde{C}^M(\tilde{\theta}^M)$. As above, investors may exhibit a time-varying DE distortion η_t such that subjective expectations follow:

$$\tilde{\mathbb{E}}_t[S_{t+v}^{M*}] = \tilde{C}_v^M(\tilde{\theta}^M) + \left[T^M(\tilde{\theta}^M)\right]^v S_t^{M*} + \left[T^M(\tilde{\theta}^M)\right]^v \zeta \eta_t \quad (23)$$

where $\tilde{C}_v^M(\tilde{\theta}^M) \equiv \tilde{C}^M + \tilde{T}^M \tilde{C}^M + \left[\tilde{T}^M\right]^2 \tilde{C}^M + \dots + \left[\tilde{T}^M\right]^{v-1} \tilde{C}^M$. The scalar parameter ζ governs the strength of the over- or underreaction to all shocks, with $\zeta > 0$, implying overreaction, and $\zeta < 0$ implying underreaction. As above, the distortion η_t follows a VAR(1) process, with an innovation that is proportional to the vector of perceived shocks $\tilde{\varepsilon}_t^M$:

$$\eta_t = \rho_\eta \tilde{T}^M \eta_{t-1} + \tilde{R}^M \tilde{Q}_{\xi_t}^M \tilde{\varepsilon}_t^M, \quad \rho_\eta \in [0, 1]. \quad (24)$$

Thus, η_t is a vector with elements comprised of unique decaying sums of multiple past perceived innovations $\{\tilde{\varepsilon}_t^M, \tilde{\varepsilon}_{t-1}^M, \tilde{\varepsilon}_{t-2}^M, \dots\}$.

The special case of rational expectations occurs when both the wedge vector w_θ and the scalar parameter ζ are zero.

Asset Pricing Dynamics The economy is populated by a continuum of identical investors who earn all income from trading in the stock market and a one-period nominal risk-free bond in zero net supply. Assets are priced by a representative investor who consumes per-capita aggregate shareholder payout, $D_t = K_t Y_t$.

The representative investor's intertemporal marginal rate of substitution in consumption

is the stochastic discount factor (SDF) with logarithm:

$$m_{t+1} = \ln(\beta_p) + \vartheta_{pt} - \gamma_{ra}(\Delta d_{t+1}). \quad (25)$$

where γ_{ra} is a curvature parameter and where the time discount factor is subject to an aggregate externality in the form of a patience shifter ϑ_{pt} that individual investors take as given.¹³ A time-varying specification for the subjective time-discount factor is essential for ensuring that investors are willing to hold the nominal bond at the interest rate set by the central bank's policy rule.

The first-order-condition for optimal holdings of the one-period nominal risk-free bond with a face value equal to one nominal unit is

$$LP_t^{-1}Q_t = \tilde{\mathbb{E}}_t[M_{t+1}\Pi_{t+1}^{-1}], \quad (26)$$

where Q_t is the nominal bond price, $\tilde{\mathbb{E}}_t$ denotes the subjective expectations of the investor, and $\Pi_{t+1} = P_{t+1}/P_t$ is the gross rate of general price inflation. We assume that investors have a time-varying preference for nominal risk-free assets over equity, accounted for by the term $LP_t > 1$ in (26), implying that the bond price Q_t is higher than it would be absent these benefits, i.e., when $LP_t = 1$. Taking logs of (26) and using the properties of conditional log-normality delivers an expression for the real interest rate as perceived by the investor:

$$i_t - \tilde{\mathbb{E}}_t[\pi_{t+1}] = -\tilde{\mathbb{E}}_t[m_{t+1}] - .5\tilde{\mathbb{V}}_t[m_{t+1} - \pi_{t+1}] - lp_t \quad (27)$$

where the nominal interest rate $i_t = -\ln(Q_t)$, $\pi_{t+1} \equiv \ln(\Pi_{t+1})$ is net inflation, $\tilde{\mathbb{V}}[\cdot]$ is the conditional variance under the subjective beliefs of the investor, and $lp_t \equiv \ln(LP_t) > 0$.

¹³This specification for ϑ_{pt} is a generalization of those considered in previous work (e.g., Ang and Piazzesi (2003); Campbell and Cochrane (1999); Lettau and Wachter (2007)). Combining (27) and (25), we see that $\vartheta_{p,t}$ is implicitly defined as

$$\vartheta_{pt}^p = -\left[i_t - \tilde{\mathbb{E}}_t[\pi_{t+1}]\right] + \tilde{\mathbb{E}}_t[\gamma_{ra}\Delta d_{t+1}] - .5\tilde{\mathbb{V}}_t[-\gamma_{ra}\Delta d_{p,t+1} - \pi_{t+1}] - lp_t - \ln(\beta_p).$$

Variation in lp_t follows an AR(1) process

$$lp_t - \bar{lp} = \rho_{lp} (lp_{t-1} - \bar{lp}) + \sigma_{lp, \xi_t} \varepsilon_{lp, t} \quad (28)$$

subject to an i.i.d. shock $\varepsilon_{lp, t} \sim \mathcal{N}(0, 1)$. Since lp_t is a component of preferences, distorted perceptions play no role in (28).

Let P_t^D denote total value of market equity. Using (8), $pd_t \equiv \ln(P_t^D/D_t)$ obeys the following approximate log Euler equation:

$$\begin{aligned} pd_t = & \kappa_{pd, 0} + \tilde{\mathbb{E}}_t [m_{t+1} + \Delta d_{t+1} + \beta pd_{t+1}] + \\ & + .5 \tilde{\mathbb{V}}_t [m_{t+1} + \Delta d_{t+1} + \beta pd_{t+1}]. \end{aligned} \quad (29)$$

Rewriting as a function of r_{t+1}^D and subtracting off (27), the log equity premium as perceived by the investor is:

$$\underbrace{\tilde{\mathbb{E}}_t [r_{t+1}^D] - (i_t - \tilde{\mathbb{E}}_t [\pi_{t+1}])}_{\text{subj. equity premium}} = \underbrace{\begin{bmatrix} -.5 \tilde{\mathbb{V}}_t [r_{t+1}^D] - \widetilde{\text{COV}}_t [m_{t+1}, r_{t+1}^D] \\ + .5 \tilde{\mathbb{V}}_t [\pi_{t+1}] - \widetilde{\text{COV}}_t [m_{t+1}, \pi_{t+1}] \end{bmatrix}}_{\text{subj. risk premium}} + \underbrace{lp_t}_{\text{liquidity Premium}}, \quad (30)$$

where $\widetilde{\text{COV}}_t[\cdot]$ is the conditional covariance under the subjective beliefs of the investor. The subjective equity premium has two components. The component labeled “subj. risk premium” is attributable to the agent’s subjective perception of the quantity of risk, which varies in the model with fluctuations in the stochastic volatilities of the macro shocks, driven by ξ_t . The term labeled “liquidity premium” comes from the time-varying preference for risk-free nominal debt over equity. It captures fluctuations in the pricing of risk due to factors not explicitly modeled, such as time variation in sentiment or implied risk aversion (e.g., from leverage constraints), flights to quality, or changes in the perceived liquidity and safety attributes of nominal risk-free assets (e.g., Krishnamurthy and Vissing-Jorgensen (2012)). We treat this risk-preference component as a latent random variable to be estimated.

Equilibrium An equilibrium is defined as a set of prices (bond prices, stock prices), macro quantities (interest rates, inflation, output growth, payout share), laws of motion, and investor beliefs such that macro dynamics in (14)-(18) and thus (19) are satisfied, asset pricing dynamics in (25)-(29) are satisfied, and investor beliefs are given by (20), (23) and (24).

Model Solution We solve the system of structural model equations that must hold in equilibrium using standard algorithms that preserve log-normality of the system. (See the Internet Appendix for details).

Let $S_t^A \equiv [m_t, pd_t, lp_t, \tilde{\mathbb{E}}_t(m_{t+1}), \tilde{\mathbb{E}}_t(pd_{t+1})]$ be a set of asset pricing state variables obeying (25)-(29), and let $S_t \equiv [S_t^M, S_t^{M*}, S_t^A, \tilde{\varepsilon}_t^M, \eta_t]'$. The solution to the complete structural model can be expressed as a MS-VAR in S_t :

$$S_t = \bar{C}(\theta_{\xi_t}, \tilde{\theta}_{\xi_t}) + \bar{T}(\theta_{\xi_t}, \tilde{\theta}_{\xi_t}) S_{t-1} + \bar{R}(\theta_{\xi_t}, \tilde{\theta}_{\xi_t}) Q_{\xi_t} \varepsilon_t, \quad (31)$$

where $\bar{C}(\cdot)$, $\bar{T}(\cdot)$, $\bar{R}(\cdot)$ are matrices of primitive parameters involving elements of θ_{ξ_t} and $\tilde{\theta}_{\xi_t}$, some of which vary with the Markov-switching variable ξ_t , and $Q_{\xi_t}(\cdot)$ is a matrix of shock volatilities that vary stochastically with ξ_t . The structural shocks are contained in $\varepsilon_t = (\varepsilon_t^M, \varepsilon_{lp,t}, \varepsilon_{v,t})'$, which stacks the primitive macro shocks ε_t^M , the liquidity premium shock $\varepsilon_{lp,t}$ (a feature of preferences), and the vintage errors $\varepsilon_{v,t}$.¹⁴

5 Estimation and Mapping to Data

State-Space Estimation and Filter The system of estimable equations is placed in state-space form by combining (31) with an observation equation taking the form

$$\begin{aligned} X_t &= D_{\xi_t,t} + Z_{\xi_t,t} S_t' + U_t v_t \\ v_t &\sim \mathcal{N}(0, I), \end{aligned} \quad (32)$$

¹⁴Neither $\tilde{\varepsilon}_t^M$ or η_t appears separately in ε_t because $\tilde{\varepsilon}_t^M = (\tilde{R}^M \tilde{Q}^M)^{-1} (S_t^{M*} - \tilde{C}^M - \tilde{T}^M S_{t-1}^{M*})$ is entirely pinned down S_t^{M*} (and thus by ε_t^M and $\varepsilon_{v,t}$), while η_t has an innovation that is proportional to $\tilde{\varepsilon}_t^M$.

where X_t denotes a vector of observable data and machine forecasts at time t , v_t is a vector of observation errors, U_t is a diagonal matrix with the standard deviations of v_t on the main diagonal, and $D_{\xi_t,t}$ and $Z_{\xi_t,t}$ are parameters that map X_t into corresponding theoretical restrictions that are functions of S_t . The parameters $Z_{\xi_t,t}$, U_t , and $D_{\xi_t,t}$ depend on t independently of ξ_t because some series in X_t are not available at all frequencies and/or over the full sample. As a result, the state-space estimation uses different measurement equations to include these series when the data are available and exclude them when they are missing.

We estimate the state-space representation with three volatility regimes (high/med/low) using Bayesian methods based on a modified version of Kim's (1994) basic filter and approximation to the likelihood for Markov-switching state space models. A random-walk Metropolis-Hastings MCMC algorithm is used to characterize uncertainty. The model parameters are estimated on mixed-frequency monthly, quarterly, and biannual data and, following BLM2, used subsequently along with high-frequency forward-looking data to conduct an event study to characterize market reactions to news. We outline this procedure below.

Priors A complete description of the priors is provided in Section 1 of the Internet Appendix.¹⁵ Here we discuss priors on parameters governing investor beliefs. For the wedge vector $w_\theta \equiv \tilde{\theta}^M - \theta^M$, we use a prior that is Normal, centered on zero, with standard deviation $\pm 5\%$ deviation from the objective parameter, i.e., $\tilde{\theta} = \theta(1 + w_\theta)$ where $w_\theta \sim \mathcal{N}(0, .05^2)$. For the parameter ζ governing the extent to which investors over- or underreact to perceived shocks, we use a prior that is Normal, centered on zero, with informative but loose tightness set to unit standard deviation to achieve modest shrinkage. Importantly, the priors for all of these parameters are symmetric, i.e., centered on zero, and are therefore without bias regarding the nature of the distortion. This is essential for our investigation because whether $\tilde{\theta} \geq \theta$ or $\zeta \geq 0$ could have important consequences for asset pricing dynamics. In both cases, our estimation treats these polar parametric possibilities as equally likely and accordingly ensures that both their sign and magnitude are approached as open empirical questions to

¹⁵Priors for most parameters are standard and specified to be loosely informative except where stronger restrictions are dictated by theory, e.g., risk aversion must be non-negative.

be investigated.

Machine Expectations We use machine forecasts of excess stock market returns, S&P 500 earnings growth, GDP growth, and inflation in our estimation. These machine forecasts map onto theoretical equations consistent with rational expectations, i.e., with wedge vector w_θ and the scalar parameter ζ both zero, and are based on forecasts of macro fundamentals obtained with forward iterations of:

$$S_t^{M*} = C^M(\theta^M) + T^M(\theta^M)S_{t-1}^{M*} + R^M(\theta^M)Q_{\xi_t}^M \varepsilon_t^{M*}, \quad (33)$$

which twists estimates of θ^M in the objective LOM (19) toward values consistent with the machine forecasts. The resulting estimator of θ^M therefore strikes a balance between providing a good description of historical data and ensuring that the parameters describing the objective data evolution are free from overfitting and look-ahead bias characteristic of a purely ex post estimation. Since we use observation errors in the estimation, there is no presumption that the machine exactly knows the parameters of the stylized model. Instead, the method imposes the less restrictive assumption that the machine forecasts provide a valuable signal of their real-time magnitudes.

Inferring Belief Reactions to News To infer how investor beliefs are affected by news, we apply the high frequency filtering algorithm developed in BLM2 for inferring revisions in investor perceptions about the current economic state in tight windows around news events. Even though investors price assets continuously, we assume that they observe monthly values for the real-time macro state vector S_t^{M*} and the corresponding volatility regime ξ_t , only at the *end* of each month. It follows that a news event arriving *within* the month can only be informative about the end-of-month values of S_t^{M*} and ξ_t , leading investors to update their beliefs over the values for these variables they expect will prevail.¹⁶ We refer to these

¹⁶Investors can observe the objective volatility regime sequence $\{\xi_t, \xi_{t-1}, \dots\}$ at the end of each t , but their perceived volatilities $\tilde{Q}_{M, \xi_t}$ may still differ from the objective Q_{M, ξ_t} .

intramonth updates in beliefs as revisions in *nowcasts*. They are equivalent to revisions in perceived shocks. We discuss the procedure for estimating the belief revisions briefly below, leaving detailed coverage of the general approach to BLM2.

Data The meta data-set used for this project consists of thousands of economic time series at mixed sampling intervals and spans the period January 1961 through December 2021. For the structural estimation, the observation vector often uses multiple noisy signals of the objective, underlying theoretical concept. In what follows, we provide a brief summary of the data and how it is used. A complete description of the data, sources, and mapping to the model is provided in the Internet Appendix.

Data used in structural estimation We estimate the model’s structural parameters on data at monthly or lower frequency sampling intervals (as available) from 2001:01-2021:12. Many series are used because they have obvious model counterparts, e.g., Gross Domestic Product (GDP) growth, Consumer Price Index (CPI) inflation, the federal funds rate (FFR), stock market returns, the S&P 500 market capitalization. We use real-time versus historical versions of these, as appropriate, for mappings onto S_t^{M*} and S_t^M . The ratios of U.S. corporate sector payout-to-GDP, S&P 500 earnings-to-GDP and S&P 500 dividends-to-GDP are all used as noisy signals on the payout share of output K_t . Investor expectations over multiple horizons are informed by (i) surveys of expectations on future stock returns from UBS/Gallup, the Michigan Survey of Consumers (SOC), the Conference Board (CB), the CFO Survey from the Richmond Federal Reserve Bank, converting firm-level earnings per share forecasts to S&P 500 forecasts by aggregating over the value-weighted firm-level forecasts and converting to growth forecasts, and the Consensus Forecasts of the S&P 500 index from Bloomberg (BBG), (ii) equity analysts’ S&P 500 one-year-ahead earnings growth expectations from IBES and Bloomberg (BBG), and the IBES long-term-growth expectations using the LTG expectation variable¹⁷ (iii) dividend growth expectations using S&P dividend

¹⁷When using the IBES long-term growth forecasts (LTG), we follow Bordalo et al. (2019) in aggregating the value-weighted firm-level long-term growth forecasts of the median analyst to obtain LTG at the S&P 500 level. Since there is ambiguity in the question framing, we treat LTG as an annual five-year forward

futures data following the procedure of Gormsen and Koijen (2020), (iv) expectations of future inflation and GDP growth from the Survey of Professional Forecasters (SPF), BBG, the Livingston (LIV) Survey (inflation only), and the Blue Chip (BC) Survey, (v) interest rate expectations using Federal Funds Futures (FFF), Eurodollar (ED) futures, both at multiple contract horizons, and the Blue Chip (BC) survey expectations of the FFR 12 months ahead.¹⁸ Data on the spread between the Baa corporate bond return and the 20-year Treasury bond return (“Baa spread” hereafter) are used as an additional noisy signal on the subjective equity premium, including the liquidity premium component, lp_t .

Our measure of corporate earnings deserves further mention. We use bottom-up measure of IBES Street Earnings for the S&P 500, which differs from GAAP earnings by excluding discontinued operations, extraordinary charges, and other non-operating items. This is important because Street Earnings—unlike GAAP earnings—are closely aligned with the target of equity analysts’ forecasts (IBES and BBG), which we use to measure subjective beliefs about equity cash-flows.¹⁹ Hillenbrand and McCarthy (2024) argue for the use of Street Earnings for measuring analysts’ expectations and show that the use of GAAP earnings can over-state the role of short-term expectations in analyst forecasts for price-earnings fluctuations.

Data used for news events and high-frequency filtering To estimate news-driven revisions in perceptions of the economic state S_t^{M*} , we use pre- and post-news event observations on a subset of the above series available at high frequency. These include tick level data on stock returns, the S&P 500 market capitalization, FFF and ED contract rates with

growth expectation, (i.e., annual earnings growth from four to five years ahead), since that delivers the lowest mean-square loss with the survey responses after considering a variety of long- and forward-horizon growth targets. Interpreting LTG as an expected annual n -year forward growth rate (rather than the expected annualized n -year growth rate) is consistent with the reference to the *next* full business cycle and moreover makes the stable median LTG forecast easier to reconcile with the volatile median one-year growth forecast.

¹⁸In principle, fed funds futures market rates may contain a risk premium that varies over time. If such variation exists, it is absorbed in the estimation by the observation error for these equations (Piazzesi and Swanson (2008)).

¹⁹According to the IBES user guide, analysts submit forecasts after backing out these transitory components, and IBES constructs the realized series to align with those forecasts. While analysts have some discretion over which items to exclude, Hillenbrand and McCarthy (2024) demonstrate that the target of these forecasts corresponds closely to earnings before special items in Compustat, suggesting that street earnings accurately reflect the measure analysts are targeting.

different expiries, daily BBG survey expectations on multiple variables, and daily data on the Baa spread.²⁰ In our analysis, the pre-event value is defined as either 10 minutes before the news event or the day prior, depending on data availability (daily versus minutely/tick level). Similarly, the post-event value is defined as either 20 minutes after the event or the following day. Our sample of news events includes (i) 1,234 macroeconomic data releases for GDP, CPI, U.S. unemployment, and U.S. payroll data spanning the period 1980:01-2021:12, (ii) 16 corporate earnings announcement days spanning 1999:03-2020:05, and (iii) 220 Federal Open Market Committee (FOMC) press releases from the Fed spanning 1994:02-2021:12. The corporate earnings news events are from Baker, Bloom, Davis and Sammon (2019) who conduct textual analyses of *Wall Street Journal* articles to identify days in which there were large jumps in the aggregate stock market attributed primarily to corporate earnings news with high confidence.²¹ We run the filter to obtain estimates of S_t and S_t^{M*} at high frequency pre- and post-news event, and at the end of every month from 2008:01-2021:12.

Data inputs for machine learning algorithm The algorithm used to produce dynamic machine expectations uses thousands of initial data inputs. Following BLM1, many of these are converted to diffusion index factors before being passed to the machine estimator. The initial data inputs include a real-time macro dataset on 92 indicators, a panel dataset of 147 monthly financial indicators, and daily “up-to-the-forecast” financial market information from five broad classes of financial assets: (i) commodities prices (ii) corporate risk variables including credit spreads (iii) equities (iv) foreign exchange, and (v) government securities. A number of other inputs are used, including consensus forecast surprises around data releases, FFF revisions, market jumps around past news events, and daily text-based factors estimated by Latent Dirichlet Allocation (LDA) analysis from around one million articles published in

²⁰For events that occur when the market is closed we use minutely data on the S&P 500 E-mini futures market.

²¹Baker et al. (2019) (BBDS) examine next-day newspaper accounts of big daily moves (“jumps”) in the stock market. Trained human readers classify the proximate cause of each jump into distinct categories and code the confidence with which the journalist advances an explanation for the jump. We are grateful to the authors of Baker et al. (2019) for providing us with their data for these event days.

the *Wall Street Journal* between January 1984 to June 2022.²²

6 Results

This section presents our estimation results. We begin with preliminary analysis of structural parameter estimates.

Parameter Estimates Table 2 reports the posterior mode values for model parameters. Where applicable, separate values are reported for estimates of parameters governing the objective macro LOM and the perceived macro LOM. Several results are worthy of highlighting.

Table 2: Parameters

	Objective	Perceived		Objective	Perceived
ψ_i	0.0098	0.0098	$\beta_{\pi,i}$	0.0016	0.0016
ψ_π	0.0066	0.0065	$\beta_{\pi,\pi}$	-0.0001	-0.0001
$\psi_{\Delta y}$	0.0002	0.0002	$\beta_{\pi,\Delta y}$	-0.0516	-0.0532
ϕ_i	0.0320	0.0320	$\beta_{\Delta y,i}$	-0.0043	-0.0043
ϕ_π	0.0001	0.0001	$\beta_{\Delta y,\pi}$	0.0007	0.0007
$\phi_{\Delta y}$	0.0006	0.0006	$\beta_{\Delta y,\Delta y}$	0.0006	0.0006
ϕ_k	0.0611	0.0611	$\beta_{k,\overline{\Delta y}}$	-7.1724	-7.0824
γ_{ra}	3.9131	—	$\rho_{k,k}$	0.9803	1.0000
ζ	1.2626	—	ρ_{lp}	0.4317	—
			ρ_η	0.9982	—

Notes: Posterior mode values of the parameters. The estimation sample spans 1961:M1-2021:M12.

ALT TEXT: Table of posterior mode estimates for objective and perceived parameters. Many parameter pairs are similar, while payout dynamics and zeta indicate belief distortion.

First, the scalar parameter ζ is estimated to be a positive value equal 1.26, consistent with *overreaction* to all perceived shocks in $\tilde{\varepsilon}_t^M$.

Second, for most parameters governing perceived macro dynamics, there is little deviation from the corresponding objective parameter value. However, there are two notable

²²The results are based a randomly selected subsample of 200,000 articles over the same period. This procedure follows Bybee et al. (2021) by estimating topic weights for individual articles to construct a time series of news attention by topic.

exceptions. (i) We find $\tilde{\rho}_{k,k} > \rho_{k,k}$, implying that news about today's payout share is over-extrapolated to future payout share movements. (ii) There are differences in the perceived and actual values of $\beta_{k,\overline{\Delta y}}$, which in both cases is negative. This parameter governs the effect of trend economic growth on the payout share. The negative values for these parameters indicate that increases in trend growth $\overline{\Delta y}_t$ drive *down* the payout share k_t .²³ Because k_t affects $\bar{k}_t$ and ultimately future k_t through (14) and (18), this implies that increases in $\overline{\Delta y}_t$ cause a long-lasting decline in k_t . Yet because $0 > \tilde{\beta}_{k,\overline{\Delta y}} > \beta_{k,\overline{\Delta y}}$, investors *underestimate* the absolute impact of $\overline{\Delta y}_t$ on k_t , implying that observed declines in k_t originating from increases in $\varepsilon_{\overline{\Delta y},t}$ are partly misattributed to another impulse that also would be perceived to cause k_t to decline. We return to this below. At the same time, $\beta_{\Delta y,\Delta y}$ is positive but small, indicating that $\overline{\Delta y}_t$ has only modest predictive power for future output growth, consistent with the fact that output growth is not highly autocorrelated. Putting this all together, increases in $\varepsilon_{\overline{\Delta y},t}$ are tantamount to bad cash-flow news: the positive effects of trend growth on future output growth are outweighed by the persistent negative effects on the payout share of output.

Third, although not shown in Table (2), Table A.7 in the Internet Appendix shows that the structural shocks with the largest standard deviations are those to the payout share of output, especially those to the cyclical payout share. This estimate is driven in the data by a highly volatile corporate earnings-to-GDP ratio. These results are consistent with the idea that, on average, news generates more uncertainty about the share of output that will ultimately accrue to profits than it does about macroeconomic aggregates or discount rates. Since the payout share is stationary, these estimates are consistent with a large negatively autocorrelated component in payout growth generated by fluctuations around a trend.

Table 3 shows basic asset pricing moments for stock returns and the real interest rate implied by these estimates. The model-implied moments for these series are based on the modal parameter and latent state estimates and match their data counterparts closely.

²³This result echoes findings in Greenwald et al. (2025), which shows that the U.S. stock market grew far faster during decades with sluggish economic growth but rapid growth in the earnings share, than in decades with rapid economic growth but a relatively stable earnings share.

Table 3: Asset Pricing Moments

Moments	Model		Data	
	Mean	StD	Mean	StD
Log Stock Return	8.75	12.32	8.96	12.29
Log Excess Return	7.27	14.82	7.42	14.85
Real Interest Rate	1.48	2.90	1.54	2.53


Notes: Model moments based on modal parameter and latent state estimates. Annualized monthly statistics (means multiplied by 12, standard deviations by $\sqrt{12}$) and reported in units of percent. The log return (data) is the log difference in the S&P 500 market cap; excess returns subtract off FFR. The real interest rate is FFR minus the average one-year ahead forecasts of inflation from the BC, SPF, SOC, and Livingston surveys. The sample is 2001:M1 - 2021:M12.

ALT TEXT: Table comparing model and data moments. Model means and standard deviations closely match the data for stock returns, excess returns, and the real interest rate.

We close this section by emphasizing that, even though the same positively estimated DE parameter ζ governs magnitude of overreactions to all shocks, the perceived volatility and propagation properties of the shocks themselves can generate asymmetries in these overreactions. It is therefore instructive to consider which shocks investors are most overreactive to as a result of this same DE parameter. Figure 1 shows how the investor's expected payout growth responds to different perceived (two standard deviation) shocks, v periods after the shock, under the estimated parameters of our baseline model. These responses are juxtaposed with those under the restricted set of parameter values corresponding to RE. A third, "No DE," line shows the responses when $\zeta = 0$ while keeping the estimated distortions between θ and $\tilde{\theta}$. A comparison of the baseline and RE responses shows that the largest such overreactions (in deviations from steady-state) are to the cyclical payout share shock first, and to the trend payout share shock second. Other shocks have smaller or negligible overreactions.²⁴ We discuss the responses to the two payout share shocks and their contribution to market volatility further below. We also see that the only shock for which there is a non-trivial difference between the baseline responses and the No-DE responses are those to the trend growth shock.

²⁴The responses to the liquidity premium display no distortion since that is driven by risk preferences and not beliefs.

Figure 1: Impulse Responses of Expected Payout Growth



Notes: This figure plots estimated impulse responses of expected payout growth at the posterior mode parameter values, in deviations from steady-state, to shocks specified in the subpanel headers.

ALT TEXT: Line graphs with nine panels showing impulse responses of expected payout growth to each modeled shock.

Market Reactions to News: High-Frequency Structural Event Study We now turn to the question of how markets react to real-world news events. To do so, we use the BLM2 filtering algorithm to infer revisions in investor perceptions about the economic state, at high frequency around news events.

This procedure can be summarized as follows. Consider news events that occur within a given month t . Let $\delta_h \in (0, 1)$ represent the number of time units that have passed during month t up to and including some particular point $t - 1 + \delta_h$. Let $S_{t|t-1+\delta_h}^{M*(i)}$ denote a filtered estimate of investor beliefs as of time $t - 1 + \delta_h$ about the time- t economic state investors expect to prevail when it is observed at the end of the month. This is an estimate of the investor's nowcast of $S_t^{M*(i)}$, conditional on the volatility regime $\xi_t = i$. Let the associated filtered volatility regime probabilities be denoted $\pi_{t|t-1+\delta_h}^i \equiv \Pr(\xi_t = i | X_{t-1+\delta_h}, X^{t-1})$, where X^{t-1} denotes the history $\{X_{t-1}, X_{t-2}, \dots\}$. Finally, let δ_h assume distinct values δ_{pre} and


δ_{post} , denoting the moments immediately before and after the news event. Announcement-specific revisions in $S^{M*(i)}$ and in π^i are computed using high-frequency, forward-looking data by taking the difference between the estimated values for these variables pre- and post-news event. These differences can be linked back to jumps in specific variables in X_t (e.g., the stock market) using the mapping (32) and further decomposed into contributions attributable to revisions in the investor's *perceived shocks* and volatility regimes using (31). We refer to these announcement-specific jump decompositions as “shock decompositions” and report them below for the stock market. To estimate the contribution of movements in subjective return premia, we report the combined contributions of lp_t and the volatility regimes to fluctuations in $\tilde{\mathbb{E}}_t[r_{t+1}^D] - (i_t - \tilde{\mathbb{E}}_t[\pi_{t+1}])$ in (30), labeled “equity premium” in the figures below.

For the macroeconomic data releases and Fed news events we have an exact time stamp indicating when the information was released to the public. This allows us to construct precise 30 minute windows for these events (δ_{pre} =10 minutes before to δ_{post} =20 minutes after). We then run the filter at these times pre- and post-news using minutely or tick-level financial market data. We also use daily data on the day before and the day after these events for those series that are available daily but not at higher frequency. For the corporate earnings news—where events span an entire day—we run the filter using information on all high-frequency series from the close of the market on the day before to the opening of the market on the day after.

Our structural event study findings are divided by news category and displayed as a series of bar charts, with the news event itself described in subpanel titles. For each shock decomposition figure, we report the stock market jump in the data (as measured by the S&P 500) with a black dot and the jump implied by the estimated model with a red triangle. For the estimated baseline model (“Base”, shown in the first bar from the left), the black dots always lie on top of the red triangles because the baseline model is able to match both the direction and magnitude of the market jump with negligible error. We then compare these decompositions to two counterfactuals that reveal how the market would have behaved in

the absence of distortions: (i) rational expectations (RE), i.e., $\zeta = 0$ and $w_\theta = 0 \forall \theta$, and (ii) No DE, i.e., only $\zeta = 0$. As we have 1,470 separate news events, for the purposes of the plots below we focus on the news by category associated with the biggest stock market jumps in absolute value. (The results for all events are summarized in a subsequent table.)

Figure 2: Decomposing Jumps in S&P 500 due to FOMC News



Notes: The figure reports shock decompositions of the pre-/post- FOMC announcement change in S&P 500 attributable to revisions in the perceived macro shocks and the subjective equity premium (the combined effect of shocks to lp_t and stochastic volatility). The specific FOMC events reported on are those coinciding with the four largest jumps in the S&P 500 in the high-frequency event window. The modifiers (+) or (-) refer to the sign of the baseline response to a positive increment in the fundamental shock labeled in the legend. The sample is 2001:M1-2021:M12.

ALT TEXT: Stacked bar charts with four panels decomposing four large FOMC related S&P 500 jumps into perceived macro shocks and subjective equity premium components.

Figure 2 reports results for Fed news events. Panel (a) depicts the market response to the most quantitatively important Fed announcement in our sample, which occurred on January 3, 2001, when the central bank announced it would decrease the target for the federal funds rate by 50 basis points, resulting in a 4.1% surge in the S&P 500 during the 30 minute window surrounding the news. Figure 2 shows what the representative investor learned from


this announcement, as seen through the lens of the model. For the baseline model (leftmost decomposition) the biggest contributors to the jump were upward revisions in the perceived shocks to the trend payout share and cyclical output growth, and a downward revision in the subjective equity premium. Under the RE counterfactual, the market would have jumped up 3.4% rather than 4.1%. This overreaction is driven almost entirely by the DE distortion, a finding that can be observed by noting that the “No DE” counterfactual results in virtually the same jump and decomposition as the RE counterfactual. DE causes investors to react with excessive optimism to both the trend payout share and cyclical output growth shocks, inflating the price response. This same pattern leads to even greater overreaction for the FOMC event depicted in panel (d), when the market jumped up by 2.10%, while it would have jumped only 1.48% under RE.

An important additional finding that can be observed from Figure 2 is that distortions are negligible in response to shocks that drive perceived risk premia. We emphasize that this is not by construction but is instead an empirical result. As Table A.7 in the Internet Appendix shows, the shock volatilities move considerably over time and the agent’s subjective perception of the quantity of risk is allowed to differ in estimation from the objective perception. Thus, in principle, the agent’s subjective risk premium could have differed from the objective risk premium, but this is not what we find.

Figure 3 analyzes market reactions to news about the macro economy. Panel (a) shows that, on March 11, 2020, early in the Covid-19 outbreak, the market fell 2.7% in the 30 minutes surrounding the Bureau of Labor Statistics (BLS) release of the CPI report, which came in .1% below consensus forecasts.²⁵ In this case the main driver of the 2.7% decline was an outsize reaction to a higher perceived cyclical payout share shock, shown in purple. Investors overreact to the expected pay-back growth induced by the higher cyclical payout share shock. However, this shock plays virtually no role in the rational response, a point we come back to below. For this reason, under the RE counterfactual, the market would have declined just 0.56%. The same pattern plays out in reverse for the event in panel (b).

²⁵The BLS releases typically occur at 8:30 am. We use S&P E-mini futures data to gauge market reactions to these events.

Figure 3: Decomposing Jumps in S&P 500 due to Macro News




Notes: See Figure 2. The figure reports shock decompositions for the four biggest macro news events based on absolute jumps in the S&P 500 in the high-frequency event window.

ALT TEXT: Stacked bar charts with four panels decomposing four large macro news S&P 500 jumps into perceived macro shocks and subjective equity premium components.

Surrounding the release of data showing surprising economic weakness, estimates imply a downward revision in the perceived cyclical payout share shock. This revision contributed strongly to the 2% jump upward in the market, as investors overreacted to expected catch-up growth. By contrast, the market would have increased only 0.42% under RE, mostly due to a declining equity premium.

To understand these results, we must first explain why a positive increment to the cyclical payout share shock causes a sharp decline in the stock market in the baseline case, but not under RE. For this we refer to Figure 4, which plots estimated impulse responses, in deviations from steady state, to 2 standard deviation increases in the cyclical payout share shock, $\varepsilon_{k,t}$, (top row) and in the trend payout share shock, $\varepsilon_{\bar{k},t}$, (bottom row). From the top row, left panel, we see that a positive payout share shock leads to a highly transitory

Figure 4: Impulse Responses to Payout Share Shocks



Notes: This figure plots estimated impulse responses at the posterior mode parameter values, in deviations from steady-state, to positive payout share shock (top panels) and a positive trend payout share shock (bottom panels), respectively.

ALT TEXT: Line graphs with six panels comparing responses to cyclical payout share and trend payout share shocks.

increase in payout relative to GDP that quickly mean reverts. Under both RE and in the baseline model, this mean reversion in the payout *share* creates the expectation of negative fall-back *growth* in payout next period (top row, middle panel), consistent with the common understanding that the payout share is stationary. However, in the RE case, expected fall-back growth is just negative enough to (almost) restore k_t to its steady state level within one period, so that expected growth from period 1 onward is approximately zero. Because the shock is rationally perceived to generate a highly transitory deviation from the steady-state payout-output share, it has a negligible effect on the level of the stock market (top row, right panel).²⁶ By contrast, in the baseline model with DE, the investor strongly

²⁶Expected growth period 1 onward would be exactly zero if the effects of $\varepsilon_{k,t}$ on payout growth were exactly i.i.d, but these shocks affect the trend $\bar{k}_t$ (see 14), albeit with an estimated positive loading that is quite small. In turn, $\bar{k}_t$ feeds back into next period's k_t (see 18). Taken together, these features imply that a positive impulse to $\varepsilon_{k,t}$ has—via its small effect on the trend payout share—a small (almost imperceptible)

overreacts to the increase in $\varepsilon_{k,t}$, giving rise to excessive pessimism about fall-back growth next period. This effect (demonstrated above in the simplified model) is tantamount to temporarily believing that the level of the payout share will revert back to permanently lower level. This erroneous belief causes the market to crash before recovering next period when actual growth is observed and investors observe that they had been excessively pessimistic (top row, right panel).

These results can be contrasted with the responses to a trend payout share shock in the bottom row of Figure 4. Unlike an increase in the cyclical payout share shock, $\varepsilon_{k,t}$, an increase in $\varepsilon_{\bar{k},t}$ has highly persistent effects on k_t , implying that mean reversion takes decades.²⁷ When the shock hits, we get a near-permanent rise in the payout share and a one-time spike up in expected payout *growth*. In the baseline model, the good news from $\varepsilon_{\bar{k},t}$ is overreacted to, creating excessive optimism and inflating the price response. The same phenomenon leads to disappointment the following period when actual growth is observed and investors learn that they had been excessively optimistic, causing a gradual price reversal back toward the RE response.²⁸

Returning to Figure 3, Panel (c) shows reactions to the June 5, 2020 BLS release of the unemployment rate, in which the stock market rose 2.01%, an increase that coincides with the baseline model response. In this case, however, the market's response to this news would have been to jump 3.47% had expectations been formed rationally, implying that distorted beliefs and DE in particular led the market to *underreact* to the June 5, 2020 announcement. This happens because, while the news causes investors to revise down their perception of the equity premium, which pushes the market up, in the baseline model it also causes a partially

yet highly persistent effect on rationally expected future payout growth starting in period 1. It is the persistent movement in $\bar{k}_t$ that causes the RE stock price *level* (as opposed to the price-payout *ratio*) to rise imperceptibly above zero on impact, before declining slowly back to baseline over time (third panel). Since the payout share mean reverts over time regardless of its persistence, positive shocks to the share would influence prices (discounted forward-looking cash flows) less than current cash flows leading to a lower price-payout ratio.

²⁷Payout rises in period 1 because $\varepsilon_{\bar{k},t}$ affects k_t with lag—see (18).


²⁸The baseline price remains slightly above the RE level for some time before finally converging. This happens because excessive optimism or pessimism generated by DE is modulated by the shock's perceived persistence, which in this case is estimated to be high—see (24).

offsetting upward revision in the cyclical payout share shock, which pushes the market down. As in panel (a), this occurs because an increase in the perceived cyclical payout share causes excessive pessimism about expected fall-back growth next period.

Under RE, there is no excessive pessimism to the cyclical payout share shock and thus no erroneous partially offsetting contribution that dampens the market response. In addition, under RE we see that a downward revision in the perceived shock to trend growth $\overline{\Delta y}_t$ makes a large positive contribution to the market change (yellow bar) because, as the parameter estimates in Table 2 indicate, a negative shock to trend growth generates an expectation of higher future payout growth. By contrast, this same perceived shock makes a smaller positive contribution to the market change in the no-DE case due to distortions in the perceived LOM that underestimate this effect, which means that changes in $\varepsilon_{\overline{\Delta y},t}$ will be partly misattributed to another shock. As Figure (6) shows, in this case investors erroneously attribute part of these effects to a positive cyclical payout share shock, which also drives up payout but less persistently than does the negative impulse to $\varepsilon_{\overline{\Delta y},t}$. This misperception therefore dampens the effect of the true $\varepsilon_{\overline{\Delta y},t}$ impulse on the market, an outcome that is necessarily amplified by DE, since DE applies to the shocks that investors perceive rather than those that actually occurred. This amplified underreaction explains why the contribution of the yellow bar in panel (c) is so small. Overall, this event illustrates the potential for DE to cause underreaction to news in a multi-shock setting. It is important to emphasize that this underreaction is not due to inattention. In this model, a single parameter with an estimated value indicative of behavioral overreaction controls the distorted reactions to all shocks. The event in panel (c) underscores the capacity of DE to generate asymmetric compositional effects capable of dampening market fluctuations in a multi-shock environment.

Figure 5 shows shock decompositions for the stock market's reaction to news about corporate earnings on big corporate earnings news days. Consider January 20, 2009, a news day in the wake of the financial crisis when the market declined 5.6% amid extensive reports about unrealized losses in the portfolio of asset manager State Street and in those of large banks—panel (b). Seen through the lens of the model, events like these can create the

Figure 5: Decomposing Jumps in S&P 500 due to Earnings News




Notes: See Figure 3. The figure reports shock decompositions for the four biggest corporate earnings news events based on absolute jumps in the S&P 500 in the high-frequency event window.

ALT TEXT: Stacked bar charts with four panels decomposing four large earnings news S&P 500 jumps into perceived macro shocks and subjective equity premium components.

expectation of lower output and a temporarily higher payout *share* of output that will shortly mean revert, creating the expectation of lower payout growth going forward. The baseline model implies that the market declined largely because investors became overly pessimistic about fall-back growth in payout, after revising their perception of the short-term, cyclical component of the payout share upward. This effect on the market was only partially offset by a downward revision in the subjective equity premium. By contrast, under RE, the market would have declined just 0.7% in response to this news day, largely because there is no overreaction to the cyclical payout share. At the same time, panel (b) of Figure 3 shows that, under RE, an upward revision in the perceived shock to trend growth $\overline{\Delta y}_t$ (yellow bar) makes a large negative contribution to the market response that is barely visible in the plot for the baseline model. This same phenomenon arises in the event of panel (c) of Figure 3

only in the opposite direction. As explained above, this happens because perceptions about how trend economic growth affects payout are distorted, which implies that impulses in $\varepsilon_{\Delta y, t}$ are partly misattributed to another shock.

Figure 6: Responses of Perceived Shocks to Actual Trend Growth Shock



Notes: This figure plots estimated impulse responses at the posterior mode parameter values, in deviations from steady-state, of perceived shocks to an actual (positive) trend growth shock.

ALT TEXT: Line graphs with eight panels showing perceived shocks after a positive actual trend growth shock.

To see which shocks this misattribution maps into, we report in Figure 6 estimated impulse responses of all perceived shocks in $\tilde{\varepsilon}_t^M$ to a 2 standard deviation increase in the actual trend growth shock $\varepsilon_{\Delta y, t}$. Under RE, only the perceived trend growth shock responds to an actual trend growth shock, as all perceptions are accurate. By contrast, in the baseline model, an increase in $\varepsilon_{\Delta y, t}$ not only causes $\tilde{\varepsilon}_{\Delta y, t}$ to increase, it also causes the perceived cyclical payout share shock, $\tilde{\varepsilon}_{k, t}$, to decrease strongly and persistently, and causes the perceived inflation shock, $\tilde{\varepsilon}_{\pi, t}$, to increase by a smaller absolute magnitude. The confounding negative effect on $\tilde{\varepsilon}_{k, t}$ of a positive impulse to $\tilde{\varepsilon}_{\Delta y, t}$ creates the false expectation of catch-up growth in payout next period. Since this false expectation has price effects that counteract those of $\tilde{\varepsilon}_{\Delta y, t}$, and are amplified by DE, objective changes in the trend component of economic growth are heavily dampened in the baseline model.

Table 4: Average Jump Differentials

All Events	Biggest Jumps	Smallest Jumps
Macro News		
-12.2%	24.2%	-24.9%
Corporate Earnings News		
14.4%	43.1%	5.4%
FOMC News		
-13.3%	12.1%	-18.3%

Notes: This table reports $(|J^{Base}| - |J^{RE}|) / |J^{Market}|$ the average difference between the pre-/post- news event jump (in absolute value) for the baseline model $|J^{Base}|$ and that for the counterfactual RE case $|J^{RE}|$ divided by the absolute market jump $|J^{Market}|$. For macro and Fed news, "Biggest" ("Smallest") refers to the top (bottom) 10% of all events based on the absolute change in the S&P 500 over the news window. For earnings news "Biggest" ("Smallest") refers to the top (bottom) 3 events.

ALT TEXT: Table of average jump differentials for macro, earnings, and FOMC news.


Table 4 summarizes the magnitude of over- or underreaction across all news event of a given type in our sample. We compute the difference between the absolute value of the pre-/post- news-event jump in the S&P 500 implied by the baseline model and the RE counterfactual, then average these differences across all events in a given category and express it as a fraction of the absolute jump in the market. Positive values for this difference indicate overreaction on average, while negative values indicate underreaction.²⁹ We repeat the computation for news that generated the "Biggest" and "Smallest" absolute jumps in the S&P 500 during the news window. For Fed and macro news (where we have hundreds of events) "Biggest" ("Smallest") refers to the top (bottom) 10% of all events based on the absolute change in the S&P 500. For the corporate earnings news events, where we have only 16 event-days, we define "Biggest" ("Smallest") as the top (bottom) 3 events according to the same criteria.

Table 4 shows that, averaged across all events, we find negative differentials in the categories of Fed and macro news, i.e., underreaction, a result driven by the smallest market events. The biggest market events are characterized by overreaction in all news categories. The largest of these is for earnings news, where the market overreacted by an average of 43%

²⁹We average across all events in which the baseline and non-distorted jumps are in the same direction. Jumps in the opposite direction happen infrequently in our sample, but also can't be readily categorized as either over- or underreaction, as opposed to simple "wrong" reaction.

of the total market change. Two points are worth noting. First, the corporate earnings news events are only large events, as BBDS focus on days with large stock market movements. Second, many of the macroeconomic data and FOMC press releases convey little if any information that was not already anticipated. Naturally, these events reside in our “Smallest” events category because they generate little or no reaction in the market and thus little or no over- or underreaction in absolute terms, even though the latter can still be large as a percentage of a tiny market change.

Figure 7: Counterfactual Simulations of S&P 500-GDP Ratio



Notes: Panel (a) plots the log S&P500-to-lagged GDP ratio, $p_t^D - y_{t-1}$ (blue solid line) along with the model-implied $p_t^D - y_{t-1}$ in dashed red line, obtained using the modal estimates of the parameters and latent states. Panel (b) plots the data for $p_t^D - y_{t-1}$ in solid blue and the counterfactual simulation of RE for $p_t^D - y_{t-1}$ in dashed blue, where both series are standardized. Panel (c) plots a historical decomposition for the GFC of changes in $p_t^D - y_{t-1}$ relative to 2007:M1. Panel (d) plots the counterfactual historical decomposition under RE of the same series. Panels (e) and (f) present the same for the Covid crash relative to 2020:M1. The black lines in (c)-(f) plot the changes in $p_t^D - y_{t-1}$ relative to date indicated in the vertical bar for each case. The sample in panels (a) and (b) spans 2001:M1 - 2021:M12.

ALT TEXT: Six-panel line and bar charts comparing the S&P 500 to lagged GDP ratio in data, the baseline model, and a rational expectations counterfactual.

Market Valuation: Historical Analysis We now study the model’s predictions outside of tight windows around news events. Panel (a) of Figure 7 reports the log ratio of market equity to last month’s output, $p_t^D - y_{t-1}$, for both the data³⁰ and the baseline model, where the latter is computed at the modal values of all parameters and latent states. (Because the baseline model effectively fits the observed series without error, two lines lie on top of each other.) Panel (b) reports the data once more, along with our estimate of the market evolution under a counterfactual simulation in which parameters that are consistent with RE prevailed. Note that a counterfactual simulation feeds in the shocks implied by the baseline model estimates while changing only the parameter values, a procedure that isolates the strength of the mechanism in the baseline estimates compared to some counterfactual mechanism.³¹ Since the counterfactual will have a different starting value by construction, we standardize both series in panel (b) to facilitate comparison. The plots in (a) and (b) span 2001:01-2021:12. The two bottom rows of Figure 7 reports historical decompositions of the variation in $p_t^D - y_{t-1}$ during two specific episodes: the Global Financial Crisis (GFC) in row 2, and the Covid-19 Crisis in row 3. These decompositions are reported for the baseline model in panels (c) and (e) and for the RE counterfactuals in panels (d) and (f). The black lines represent the cumulative month-to-month changes in $p_t^D - y_{t-1}$ relative to a start date for the episode, while the colored bars decompose these changes into cumulative fundamental shocks and premia.

While panel (a) says that the baseline model perfectly explains the market’s fluctuations, panel (b) tells us that the fit of the counterfactual RE case is far worse. A counterfactually rational stock market would have been more volatile than actually observed, resulting in a puzzle of “excess stability” rather than excess volatility. This finding demonstrates the extent to which distorted beliefs with behavioral overreaction were a stabilizing force over

³⁰We use the interpolation method of (Stock and Watson (2010)) to obtain a monthly GDP series for estimation.

³¹This differs from the implications of a counterfactual *model*, in which the shocks are re-estimated under an alternative set of parameters not chosen by the baseline estimation. The latter may be of interest in some contexts, but it cannot isolate the strength of a baseline model mechanism, since both the mechanism and the shocks change.

the post-millennial period, substantially cushioning declines during the GFC and the Covid crisis, among other episodes.

The historical decompositions in the bottom row help to explain this result. The GFC episode is characterized by a sharp decline in the cyclical payout share, to which the investor strongly overreacts, leading to excessive optimism about catch-up growth in payouts. That over-optimism makes a positive contribution to the market (purple bar), partially offsetting the predominating negative contributions due to other shocks that were still overreacted to but to a lesser degree. The market declines more under the RE counterfactual because, in that case, there is no overreaction to the decline in $\varepsilon_{k,t}$ and thus no excessive optimism about catch-up growth. This underreaction in the GFC is a prominent historical example of the shock composition effect at work. Similarly, the third row shows that, following the outbreak of the Covid-19 pandemic, the RE counterfactual simulation of $p_t^D - y_{t-1}$ again declines by more than the baseline series, though the difference is smaller. In this case, both over-optimism about catch-up growth in payouts and overreaction to a perceived decline in the trend component of real interest rates (a positive for the market) generate a smaller market decline relative to the RE counterfactual where neither over-optimism nor overreaction are operative.

7 Unpacking the Mechanism

To unpack the model's main mechanisms on belief reaction to news, this section presents several results that illustrate how markets can underreact to news even when investors overreact to all shocks.


It is instructive to begin by examining the events that produced the largest estimated underreactions, which exhibit patterns representative of most underreaction events. Panels (a)-(d) of 8 plot shock decompositions for the four events in our sample that produce the largest underreactions. Three points deserve emphasis. First, in each case, the investor in our model perceives good discount rate news simultaneously with bad cash-flow news. And,

in each case, the underreactions are attributable to asymmetries in the distorted reactions to counteracting fundamental shocks (the shock composition effect). The market rises “too little” because the investor’s expectations for earnings growth are more overly pessimistic than her views on discount rates are overly rosy. Second, the figure shows which components of the underlying data are at work to generate this finding. Specifically, many of these events occur in periods of economic weakness when the Fed was cutting interest rates while the outlook for earnings growth was bleak. Panels (e)-(g) of Figure 8 show the movements in the high-frequency data that drive these model estimates. Around all of these events, the BAA spread jumps down, leading us to estimate a decline in the equity premium. At the same time, the daily BBG earnings nowcast (relative to GDP) jumps up, while the BBG 1yr earnings growth forecast jumps down. The two combined illustrate why, as shown in panel (h), we estimate a strong upward revision in the investor’s expectation for the cyclical component of the payout share resulting in lower future cash-flow growth due to pay-back, to which the investor strongly overreacts.³² Third, at the same time, panel (h) also implies that excessive pessimism about future cash-flow growth surrounding these events is not solely attributable to the cyclical component of payout growth. Investors are also excessively pessimistic about the trend component of payout growth, a finding exhibited in panels (a)-(d) by the larger contribution revisions in these perceived shocks make to the baseline market reaction as compared to the RE counterfactual reaction. Still, the cyclical payout share shock makes a larger contribution to the underreaction than the trend component does, as exhibited by the fact that the purple bars in the baseline case are larger than the magenta.

These results suggest that the cyclical payout share shock plays an important role in our finding that overreaction to all shocks can be a force for stability rather than volatility. To quantify its importance, Figure 9 reports the results of a counterfactual simulation in

³²We emphasize that these results are not a mechanical result of our estimate that the wedge between the subjective and objective equity premia is small. Indeed a different and opposite result could have arisen if any of these had been true: (i) investors overreact to only a single cash-flow shock, rather than multiple primitive shocks with separate high- and low-frequency components (the structure matters); (ii) the events themselves did not generate jumps in multiple high-frequency variables with counteracting implications for valuations (the data matters), and/or (iii) the magnitude of the estimated overreactions to the cash-flow shocks was found to be small/negligible (the estimates matter).

Figure 8: Largest Underreaction Events (%)




Notes: The panels (a) - (d) of the figure report shock decompositions of pre-/post- FOMC announcement changes in S&P 500 attributable to revisions in the perceived macro shocks and the subjective equity premium (the combined effect of shocks to lp_t and stochastic volatility). The specific FOMC events reported on are those for which the absolute difference between the market's jump under the RE counterfactual and the baseline model as a fraction of the market jump is largest. The panels (e) - (h) show jumps in the data for these events. Panel (g) reports the BBG 1-year earnings growth forecast and panel (h) reports the estimated investor 1-year payout growth forecast from the model. The modifiers (+) or (-) refer to the sign of the baseline response to a positive increment in the fundamental shock labeled in the legend. The sample is 2001:M1-2021:M12.

ALT TEXT: Eight-panel bar charts decomposing the largest FOMC underreaction events and showing associated market, dividend, earnings, and payout growth jumps.

which the realized cyclical payout share shock is set zero. Panels (a) and (b) zoom in on the GFC. To form a basis for comparison, Panel (a) displays the log ratio of market equity to last month's output, $p_t^D - y_{t-1}$, in both the baseline model and the RE (i.e., $\zeta = 0$ and $w_\theta = 0 \forall \theta$) and No DE (i.e., $\zeta = 0$) counterfactuals *without* shutting down the cyclical payout share shock. (The baseline model and data line lie on top of each other in solid blue.) These results may be compared with Panel (b), which plots the same three cases,

but now counterfactually setting the realized cyclical payout share shock to zero. We see that, for the GFC period, counterfactually eliminating the cyclical decline in the payout share that occurred during this time—and along with it the investor’s excessive optimism about next period’s catch-up growth—would lead to the traditional finding that overreaction generates excess volatility. This underscores the central role played by assumptions on cash-flow dynamics for results on over- and underreaction. At the same time, it is important to bear in mind that the trend-cycle specification that we estimate is strongly preferred by the data, as shown above.

Figure 9: Counterfactual Simulations of S&P 500-GDP Ratio: Base vs RE



Notes: Panel (a) and (b) plots the model base estimate of $p_t^D - y_{t-1}$ in solid blue, and the counterfactual simulation of RE for $p_t^D - y_{t-1}$ in dashed blue for the GFC of changes in $p_t^D - y_{t-1}$ relative to 2007:M1. Panel (c) and (d) plot the base estimate for $p_t^D - y_{t-1}$ in solid blue and the counterfactual simulation of RE for $p_t^D - y_{t-1}$ in dashed blue, where both series are standardized. Panel (b) and (d) plot the counterfactual where the cyclical PO share shock has zero variance. The sample spans 2001:M1 - 2021:M12.

ALT TEXT: Four-panel line charts comparing baseline and rational expectations simulations of the S&P 500 to lagged GDP ratio, including simulations without cyclical payout share shock.

However, we know from Figure 8 that asymmetries in the distorted reactions to counteracting fundamental shocks (the shock composition effect) can generate excess stability

even without the cyclical shock. This is illustrated in panels (c) and (d) of Figure 9 which analyzes the full sample rather than just the GFC. For ease of reference, Panel (c) reproduces panel (b) of Figure 7, which shows our baseline model's full sample implication for the market equity-output ratio, $p_t^D - y_{t-1}$, and the corresponding RE counterfactual. Panel (d) shows the full-sample alternative counterfactual (solid line) that sets the variance of the cyclical payout share shock to zero and is juxtaposed with the RE counterpart of this no-cyclical-payout counterfactual. In this case, our finding that a counterfactually rational stock market would have been more volatile than actually observed reemerges, though it is less pronounced than in the baseline model that features a highly transitory component in cash-flow growth driven by empirically relevant variation in the earnings share of output. Taken together, panels (b) and (d) demonstrate the key quantitative role of the cyclical component in the payout share for the overall magnitude of our excess stability finding, while also showing that the key mechanism holds even without these shocks due to asymmetries in the distorted reactions to counteracting fundamental shocks (the shock composition effect).

To provide a sense of the importance of the DE distortion compared to distorted perceptions about the macro LOM, Figure 9 also shows a red dashed-dotted line displaying results for a counterfactual simulation of the No-DE case i.e., with $\zeta = 0$ but keeping estimated distortions on the perceived parameters of the macro LOM. We find that distorted perceptions about the LOM driving fundamentals alone play a modest but non-trivial role in the baseline model output, with the No-DE case lying in between the RE counterfactual (but closer to it) and the baseline output. Comparing the red and blue (base) lines isolates the marginal effect of DE and again shows that it is the most important distortion we estimate.³³

As additional evidence of a strong cyclical component in investor beliefs, we examine the estimated cyclical payout share shock and its relation to actual survey forecast errors. Our

³³The finding that DE contributes to excess stability in a multi-shock context is robust to varying ζ within empirically plausible ranges that differ from the parameter mode estimate. It holds, for example, when $\zeta = 0.91$, corresponding to the *a-posteriori* 1st percentile. In that case, the standard deviation of $p_t^D - y_{t-1}$ is 0.29 during the GFC (compared with 0.24 under the baseline and 0.40 under the *No DE counterfactual*) and 0.30 over the whole sample (compared with 0.25 under the baseline and 0.40 under the *No DE counterfactual*).

procedure produces a filtered estimate of this cyclical (i.e., short-run) shock, ε_{kt} , which we observe making large contributions to investor beliefs and stock market fluctuations during the GFC in panel (c) of Figure 7, as well as around other news events in Figures 2, 3, and 5. The model estimates imply that investors overreact to downward (upward) impulses in this shock with over-optimism (-pessimism) about catch-up (fall-back) growth. Thus, a *negative* (positive) impulse in ε_{kt} should be associated with excessively positive (negative) expectations of future cash-flow growth when compared with actual future outcomes, implying a *positive* relation between ε_{kt} and survey forecast errors when measured as the actual minus forecasted value. To check this, we regress IBES or consensus (Bloomberg) forecast errors for future S&P 500 earnings growth, $\Delta e_{t+v} - \mathbb{F}_t[\Delta e_{t+h}]$ over various future horizons h , on our estimates of ε_{kt} . Table 5 shows that the coefficient from such a regression using monthly data from 1990:M1 to 2021:M12 is positive and strongly statistically significant, consistent with the predictions of the model. This shows that the model implications for overreactions to measured payout news, summarized by jumps in ε_{kt} , reflect actual overreactions in earnings surveys.

Table 5: Earnings Forecast Errors and Earnings Share Shock

Regression: $\Delta e_{t+v} - \mathbb{F}_t[\Delta e_{t+v}] = a_v + b_v \varepsilon_{kt} + \epsilon_v$					
Survey IBES: 1981:M12 to 2021:M12					
v (months)	3M	6M	9M	12M	24M
b	5.83***	2.34***	1.25***	0.76**	-0.14
	(7.68)	(4.35)	(2.96)	(2.15)	(-1.13)
R^2	0.31	0.18	0.11	0.08	0.01
obs.	250	247	244	241	229
Survey BBG: 1990:M1 to 2021:M12					
v (months)	3M	6M	9M	12M	
b	6.28***	2.60***	1.43***	0.92**	
	(9.22)	(5.22)	(3.61)	(2.46)	
R^2	0.36	0.22	0.15	0.11	
obs.	250	247	244	238	

Notes: The table reports the OLS coefficient, heteroskedasticity and serial correlation robust t -statistics (in parentheses), R^2 statistics, and number of observations from monthly regressions of v -month-ahead forecast errors of earnings growth on the cyclical payout share shock. Regressions using the IBES survey span the period 1981:M12 to 2021:M12. Regressions using Bloomberg (BBG) span 1990:M1 to 2021:M12.

ALT TEXT: Table of earnings forecast error regressions using IBES survey in the top panel, and BBG survey in the bottom panel.

To close this section, we discuss two additional checks, the output for which is placed in the Internet Appendix to conserve space.

First, for all high-frequency news events, we use the structural model to categorize the event by whether the market exhibited over- or underreaction. This is accomplished by comparing the actual market reactions with the model's RE counterfactual reaction. We then ask whether the observed high-frequency jumps in the market that—according to the model—can be distinguished as either over- or underreactions predict future returns with the appropriate sign. Table A.6 in the Internet Appendix shows that upward (downward) jumps in the market around measured news events that the model categorized as overreactions are followed by lower (higher) subsequent returns. By contrast, those categorized as underreactions have the opposite pattern, with jumps upward (downward) followed by higher (lower) future returns. The table reports variation in the precision of these estimates, with statistical significance at the 10% level generally holding, and at the 5% level for overreaction events where the market jumped down. This result is worth remarking on since we measure

only a handful of events in any given month. Given large high-frequency variation in the stock market, predictable changes in future returns in response to these events could easily have been swamped by untracked events or noise. That the signal from this small number of events is evidently high enough to observe future reversals or momentum (as appropriate) provides supporting evidence for the existence of belief overreactions to our news events, especially negative news.

Second, Figure A.1 in the Internet Appendix shows that the model produces realistic implications for survey expectations of stock returns over time. We find the model estimates of the investor's subjective expectations are close to the survey expectations and move by the right magnitudes in times of important economic change. In particular, we observe that both survey and model expectations for returns at a one-year horizon rise sharply in the GFC, consistent with substantial increases in subjective risk premia during that episode.

8 Conclusion

We measure the nature and severity of a variety of belief distortions in market reactions to hundreds of economic news events. To do so, we use a new methodology that combines estimation of a structural asset pricing model with algorithmic machine learning to quantify bias. The structural model allows for the perceived law of motion driving macroeconomic dynamics to differ from the actual law of motion, and nests specific belief formation frameworks that include diagnostic expectations and inattention. Unlike the traditional specification of these frameworks, we allow investors to react to multiple perceived fundamental shocks, with a single estimated scalar parameter ζ controlling reactions to all shocks. We show that in this multi-shock environment, investor overreaction to all shocks can cause the market to underreact to news, dampening rather than amplifying volatility.

This theoretical possibility turns out to be empirically relevant. Our point estimates of ζ imply that investors overreact to all shocks rather than being inattentive to them. Yet we find that behavioral overreaction has been a force for market stability in the post-millennial

period. This surprising result is attributable to asymmetries in the distorted reactions to counteracting fundamental shocks, something we refer to as the shock composition effect. We show that this shock composition effect well describes the stock market's behavior in several major episodes of post-millennial history, most notably the Global Financial Crisis, in which behavioral overreaction was a force for stability rather than volatility.

A transformative idea of 20th century economic thought is that financial markets are “excessively volatile” vis-a-vis predictions of canonical theory, in which stock prices are the rational expectation of future cash-flow fundamentals discounted at a constant rate (Shiller 1981, 2000). We find that a counterfactually rational stock market would have been more volatile than actually observed, resulting in a puzzle of excess stability rather than excess volatility. By contrast, a macro-dynamic model with belief overreaction in the spirit of diagnostic expectations can perfectly explain the data, not because it amplifies volatility, but because it dampens it.

References

- ADAM, K., AND S. NAGEL (2023): “Expectations data in asset pricing,” in *Handbook of Economic Expectations*: Elsevier, 477–506.
- AFROUZI, H., S. Y. KWON, A. LANDIER, Y. MA, AND D. THESMAR (2023): “Overreaction in expectations: Evidence and theory,” *The Quarterly Journal of Economics*, 138, 1713–1764.
- ANG, A., AND M. PIAZZESI (2003): “A No-Arbitrage Vector Autoregression of Term Structure Dynamics With Macroeconomic and Latent Variables,” *Journal of Monetary Economics*, 50, 745–787.
- BAKER, S., N. BLOOM, S. J. DAVIS, AND M. SAMMON (2019): “What triggers stock market jumps?”, Unpublilshed manuscript, Stanford University.

- BARBERIS, N., A. SHLEIFER, AND R. W. VISHNY (1998): “A Model of Investor Sentiment,” *Journal of Financial Economics*, 49, 307–43.
- BASTIANELLO, F., AND P. FONTANIER (2022): “Expectations and Learning from Prices.”
- BIANCHI, F., D. Q. LEE, S. C. LUDVIGSON, AND S. MA (2025): “The Prestakes of Stock Market Investing,” Technical report, National Bureau of Economic Research.
- BIANCHI, F., S. C. LUDVIGSON, AND S. MA (2022a): “Belief Distortions and macroeconomic fluctuations,” *American Economic Review*, 112, 2269–2315.
- (2022b): “A Structural Approach to High-Frequency Event Studies: The Fed and Markets as Case History,” Technical report, National Bureau of Economic Research, No. w30072.
- BORDALO, P., N. GENNAIOLI, R. LA PORTA, AND A. SHLEIFER (2019): “Diagnostic expectations and stock returns,” *The Journal of Finance*, 74, 2839–2874.
- (2024): “Belief overreaction and stock market puzzles,” *Journal of Political Economy*, 132, 1450–1484.
- BORDALO, P., N. GENNAIOLI, Y. MA, AND A. SHLEIFER (2020): “Overreaction in macroeconomic expectations,” *American Economic Review*, 110, 2748–2782.
- BORDALO, P., N. GENNAIOLI, AND A. SHLEIFER (2018): “Diagnostic expectations and credit cycles,” *The Journal of Finance*, 73, 199–227.
- BYBEE, L., B. T. KELLY, A. MANELA, AND D. XIU (2021): “Business news and business cycles,” Technical report, National Bureau of Economic Research.
- CAMPBELL, J. Y., AND J. H. COCHRANE (1999): “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior,” *Journal of Political Economy*, 107, 205–251.

- CAMPBELL, J. Y., A. W. LO, AND C. MACKINLAY (1997): *The Econometrics of Financial Markets*, Princeton, NJ: Princeton University Press.
- CAMPBELL, J. Y., AND R. J. SHILLER (1989): “The dividend-price ratio and expectations of future dividends and discount factors,” *Review of Financial Studies*, 1, 195–228.
- CHAUDHRY, A. (2023): *The Impact of Prices on Analyst Cash Flow Expectations*, Ph.D. dissertation, The Ohio State University.
- CHEN, L., Z. DA, AND X. ZHAO (2013): “What drives stock price movements?” *The Review of Financial Studies*, 26, 841–876.
- COCHRANE, J. H. (2005): *Asset Pricing*: Revised Edition, Princeton, NJ: Princeton University Press.
- CONG, L. W., K. TANG, J. WANG, AND Y. ZHANG (2021): “Deep sequence modeling: Development and applications in asset pricing,” *The Journal of Financial Data Science*, 3, 28–42.
- DE LA O, R., AND S. MYERS (2021): “Subjective cash flow and discount rate expectations,” *The Journal of Finance*, 76, 1339–1387.
- (2023): “Which Subjective Expectations Explain Asset Prices?”, Manuscript, Wharton.
- FARMER, L. E., E. NAKAMURA, AND J. STEINSSON (2024): “Learning about the long run,” *Journal of Political Economy*, 132, 3334–3377.
- GABAIX, X. (2019): “Behavioral Inattention,” in *Handbook of Behavioral Economics*, Vol. 2 ed. by Bernheim, D., DellaVigna, S., and Laibson, D. New York, NY: Elsevier, 261–343.
- GABAIX, X., AND R. S. KOIJEN (2021): “In search of the origins of financial fluctuations: The inelastic markets hypothesis,” Technical report, National Bureau of Economic Research.

- GORMSEN, N. J., AND R. S. KOIJEN (2020): “Coronavirus: Impact on stock prices and growth expectations,” *The Review of Asset Pricing Studies*, 10, 574–597.
- GREENWALD, D. L., M. LETTAU, AND S. C. LUDVIGSON (2025): “How the wealth was won: Factor shares as market fundamentals,” *Journal of Political Economy*, 133, 1083–1132.
- GU, S., B. KELLY, AND D. XIU (2020): “Empirical asset pricing via machine learning,” *The Review of Financial Studies*, 33, 2223–2273.
- HARTZMARK, S. M., AND D. H. SOLOMON (2022): “Predictable price pressure,” Technical report, National Bureau of Economic Research.
- HECHT-NIELSEN, R. (1987): “Kolmogorov’s mapping neural network existence theorem,” in *Proceedings of the international conference on Neural Networks* Volume 3, 11–14, IEEE press New York, NY, USA.
- HILLENBRAND, S., AND O. MCCARTHY (2021): “Heterogeneous investors and stock market fluctuations,” *Available at SSRN 3944887*.
- (2024): “Street Earnings: Implications for Asset Pricing,” *Available at SSRN 4892475*.
- ISMAILOV, V. E. (2023): “A three layer neural network can represent any multivariate function,” *Journal of Mathematical Analysis and Applications*, 523, 127096.
- JIN, L. J., AND J. LI (2023): “Overreaction of Long-Term Cash Flow Expectations to Past Price Movements,” Working paper, Cornell University.
- KIM, C.-J. (1994): “Dynamic Linear Models with Markov-Switching,” *Journal of Econometrics*, 60, 1–22.
- KOHLHAS, A. N., AND A. WALTHER (2021): “Asymmetric attention,” *American Economic Review*, 111, 2879–2925.

- KRISHNAMURTHY, A., AND A. VISSING-JORGENSEN (2012): “The aggregate demand for Treasury debt,” *Journal of Political Economy*, 120, 233–267.
- LETTAU, M., AND J. WACHTER (2007): “The Term Structures of Equity and Interest Rates,” Unpublished paper.
- MANKIW, N. G., AND R. REIS (2002): “Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve,” *The Quarterly Journal of Economics*, 117, 1295–1328.
- NAGEL, S., AND Z. XU (2022): “Asset pricing with fading memory,” *The Review of Financial Studies*, 35, 2190–2245.
- PIAZZESI, M., AND E. SWANSON (2008): “Futures Prices as Risk-Adjusted Forecasts of Monetary Policy,” *Journal of Monetary Economics*, 55, 677–691.
- SHILLER, R. J. (1981): “Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?” *American Economic Review*, 71, 421–436.
- (2000): *Irrational Exuberance*, Princeton, NJ: Princeton University Press.
- SIMS, C. A. (2003): “Implications of rational inattention,” *Journal of monetary Economics*, 50, 665–690.
- STOCK, J. H., AND M. W. WATSON (2010): “Research Memorandum,” https://www.princeton.edu/~mwatson/mgdp_gdi/Monthly_GDP_GDI_Sept20.pdf.
- WOODFORD, M. (2002): “Imperfect Common Knowledge and the Effects of Monetary Policy,” in *In Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps* ed. by Aghion, P., Frydman, R., Stiglitz, J., and Woodford, M. Cambridge MA: Princeton University Press, 25–58.