Why Does Return Predictability Concentrate in Bad Times?

Julien Cujean and Michael Hasler*

ABSTRACT

We build an equilibrium model to explain why stock return predictability concentrates in bad times. The key feature is that investors use different forecasting models, and hence assess uncertainty differently. As economic conditions deteriorate, uncertainty rises and investors' opinions polarize. Disagreement thus spikes in bad times, causing returns to react to past news. This phenomenon creates a positive relation between disagreement and future returns. It also generates time-series momentum, which strengthens in bad times, increases with disagreement, and crashes after sharp market rebounds. We provide empirical support for these new predictions.

JEL Classification: D51, D83, G12, G14.

---

*Cujean is with University of Maryland, Robert H. Smith School of Business, and Hasler is with University of Toronto. We are particularly grateful to Kenneth Singleton and two anonymous referees for their insightful suggestions and comments. We would like to thank Francesco d'Acunto, Pat Akey, Maria Cecilia Bustamante, Wen Chen, Hui Chen, Peter Christoffersen, Pierre Collin-Dufresne, Michel Dubois, Darrell Duffie, Bernard Dumas, Laurent Frésard, Steve Heston, Julien Hugonnier, Alexandre Jeanneret, Yoontae Jeon, Scott Joslin, Andrew Karolyi, Leonid Kogan, Alex Kostakis, Jan-Peter Kulak, Pete Kyle, Jeongmin Lee, Danmo Li, Mark Loewenstein, Semyon Malamud, Erwan Morellec, Antonio Mele, Yoshio Nozawa, Chayawat Ornthanalai, Lubos Pastor, Lasse Pedersen, Rémy Praz, Marcel Rindisbacher, Alberto Rossi, René Stulz, Ngoc-Khanh Tran, Adrien Verdelhan, Pietro Veronesi, Jason Wei, Liyan Yang, and conference/seminar participants at the 4th Financial Risks International Forum, Collegio Carlo Alberto, EFA 2015, Eurofidai 2014, FIRS 2015, Goethe University Frankfurt, HEC Montréal, SFI-NCCR Workshop, UBC Winter Finance Conference 2015, University of California at San Diego, University of Geneva, University of Maryland, University of Neuchâtel, University of Toronto, and University of Virginia for their discussions and comments. Financial support from the University of Maryland, the University of Toronto, and the Connaught New Researcher Award is gratefully acknowledged. A previous version of this paper circulated under the title "Time Series Predictability, Investors' Disagreement, and Economic Conditions." Julien Cujean and Michael Hasler certify that they have nothing to disclose and have no conflict of interests. Corresponding author: Julien Cujean, University of Maryland, Robert H. Smith School of Business, 4466 Van Munching Hall, College Park, MD 20742, USA; +1 (301) 405 7707; jcujean@rhsmith.umd.edu; www.juliencujean.com.

From the perspective of the efficient market hypothesis, one of the most important paradigms in finance, evidence of stock return predictability is a puzzle. Leading theories explain how investors may form beliefs that lead to deviations from efficiency, focusing on *unconditional* patterns in stock returns.[1] However, mounting evidence shows that return predictability fluctuates over the business cycle, with various forms of return predictability stronger in bad times.[2] Garcia (2013), for instance, shows that news content (Tetlock (2007)) better predicts future returns in recessions. This evidence suggests that a theory of predictability must also explain how fluctuations in investors' beliefs lead observable predictable patterns to concentrate in bad times. Our objective is to provide a simple learning mechanism that gives rise to this concentration.

The main idea of the paper is that agents forecast fundamentals using different models. Some agents see the state of the economy as continuously evolving—shades of grey—throughout good and bad times, while other agents view the economy in discrete terms—black and white—with good and bad times alternating over the business cycle. These two views have been adopted as canonical frameworks of economic forecasting. The first modeling approach assumes that economic variables are continuously (linearly)

[1]Hong and Stein (1999) explain price underreaction, Daniel, Hirshleifer, and Subrahmanyam (1998) explain price overreaction, and Barberis, Shleifer, and Vishny (1998) provide conditions that may enforce both. Our goal is to develop a theory in which these phenomena alternate over the business cycle.

[2]Rapach, Strauss, and Zhou (2010), Henkel, Martin, and Nardari (2011), Dangl and Halling (2012), and Piatti and Trojani (2015) find that macro variables, such as the price-dividend ratio, have better predictive power in recessions. Similarly, Cen, Wei, and Yang (2017) and Loh and Stulz (2014) find that return predictability using investor disagreement, as proxied by the dispersion in analysts' forecasts (Diether, Malloy, and Scherbina (2002)), concentrates in recessions. We provide evidence that current excess returns and disagreement better predict aggregate future excess returns in recessions (Section IV).

related to their lagged values (Sims (1980)). Active fund managers who build forecasting models typically use this continuous-state, autoregressive specification.[3] The alternative approach assumes that economic series shift discretely (nonlinearly) across different regimes of the economy (Hamilton (1989)). Federal Reserve banks widely use this discrete-state specification to forecast business cycle turning points.[4]

This simple difference in the way agents learn about the state of the economy can explain a rich set of facts about stock returns. Specifically, there exists a positive relation between agents' disagreement and future returns. More importantly, this relation becomes stronger as economic conditions deteriorate and gives rise to positive serial correlation in returns at short horizons, a phenomenon known as "time-series momentum" (Moskowitz, Ooi, and Pedersen (2012)). This phenomenon strengthens in bad times, increases with disagreement, crashes after sharp market rebounds, and reverts over longer horizons. Overall, both disagreement and lagged returns predict future returns, and their predictive power is stronger in bad times.

The model mechanism operates as follows. All agents use the same flow of news to estimate the state of the economy, which some view as discrete and others as continuous. As a result, agents assess uncertainty differently. Those that employ a continuous-state model need not revise the uncertainty of their estimates, whereas the others must continuously reassess the uncertainty of each discrete state after they observe new information. It follows

[3] See, for instance, Grinold and Kahn (1999), chapter 10 or Stewart, Piros, and Heisler (2010), chapter 5.

[4] See, for instance, `https://fred.stlouisfed.org/series/RECPROUSM156N` and `https://fred.stlouisfed.org/series/JHGDPBRINDX`.

that agents interpret same news differently depending on economic conditions. Agents who revise their assessment of uncertainty conclude that it is countercyclical (Veronesi (1999)) and thus give little weight to news in bad times. After a long streak of bad news, good news is never "good enough" for these agents to abandon their pessimistic views. Because other agents always interpret good news favorably, opinions polarize in bad times and disagreement spikes. Variation in disagreement commands a risk premium. Since disagreement spikes are persistent, they create positive serial correlation in returns at short horizons. That disagreement spikes are countercyclical makes this phenomenon stronger in bad times.

We analyze this mechanism in a dynamic general equilibrium model populated with two agents, A and B. Agents trade a stock and a riskless bond and consume the dividends that the stock pays out. They do not observe the expected growth rate of dividends, which we call the fundamental, and use different models for the empirical process that governs it. In particular, Agent A uses a continuous-state model in which the fundamental oscillates throughout good and bad times, while Agent B uses a discrete-state model in which the fundamental alternates between a high state and a low state. Unlike Agent A, Agent B must continuously reassess uncertainty as information flows from dividends.

Because agents assess uncertainty differently, they revise their expectations at different speeds and in different directions depending on the state of the economy.[5] In particular, agents' disagreement spikes when the differ-

[5]Chalkley and Lee (1998), Veldkamp (2005), and Van Nieuwerburgh and Veldkamp (2006) obtain similar learning asymmetries. In their setup, however, the information flow fluctuates with economic conditions.

ence in their assessment of uncertainty is large. Estimating the parameters of the model, we find that dividend volatility is low, so that Agent A faces low uncertainty, while the distance between the high and low states of Agent B's model is large, so that she faces high overall uncertainty. In this context, Agent A always interprets good news favorably, but Agent B may or may not do so. Good news in good states causes Agent B to over-react, revising her expectations upwards significantly faster than Agent A. In contrast, good news in bad states is never good enough to convince Agent B that economic conditions are improving. Both the difference in adjustment speeds in good states and the polarization of beliefs in bad states generate disagreement spikes.

That Agent B faces countercyclical uncertainty causes disagreement spikes to be countercyclical. Since good states are empirically more persistent than bad states, the high and low states of Agent B's model are, respectively, close to and far from the mean of Agent A's model. This asymmetry accentuates the polarization of beliefs in bad states and dampens the difference in adjustment speeds in good states. In good times both agents face low uncertainty, they adjust their beliefs similarly, and disagreement exhibits little variation. In normal times uncertainty rises and agents' expectations adjust at different speeds, which exacerbates disagreement. In bad times agents' opinions polarize and disagreement spikes.

In our model countercyclical spikes in disagreement lead return predictability to concentrate in bad times. Because variation in disagreement—positive or negative—commands a risk premium, contemporaneous returns decrease with the *square* of disagreement. The sign of disagreement determines only the direction in which returns move in the future. News in bad times polarizes opinions—Agent B underreacts relative to Agent A—inducing future returns to move opposite to news and to persistently underreact. News in normal times precipitates a revision in Agent B's expectations—she overreacts relative to Agent A—causing future returns to persistently overreact. In good times beliefs move together and returns adjust immediately. Since disagreement spikes raise the risk premium persistently, both under- and overreaction create time series momentum. That disagreement spikes are countercyclical makes time-series momentum stronger in bad times.

The model can explain several features associated with time-series momentum. Excess returns exhibit momentum over a one-year horizon and then revert over subsequent horizons. This later phase of reversal arises through the long-run behavior of agents' consumption shares, which adjust to gradually dampen the effect of short-term disagreement spikes. Moreover, the serial correlation of returns has a hump-shaped term structure: momentum is strong at short horizons and decays at intermediate horizons. By contrast, in good times excess returns exhibit strong reversal at short horizons. An important consequence of this reversal spike is that a time-series momentum strategy may crash after sharp market rebounds.

We provide empirical support for three predictions of the model. We construct an empirical proxy for the square of disagreement using the dispersion in analysts' forecasts. Based on this proxy, the model predicts that *(1)* future excess returns are positively related to dispersion, and time-series momentum at short horizons *(2)* increases with dispersion and *(3)* is therefore strongest in bad times. We test these predictions and find that dispersion positively predicts future excess returns on the S&P 500 and that time-series momentum increases significantly with dispersion at short horizons. We show that, over the last century, time-series momentum at a one-month lag is significantly stronger during NBER recessions.

This paper contributes to the vast literature on heterogeneous beliefs. While we borrow the methodologies of David (2008) and Dumas, Kurshev, and Uppal (2009), we introduce a new form of disagreement. In this paper investors disagree because they use different forecasting models, namely, an Ornstein-Uhlenbeck process (Agent A) and a two-state Markov chain (Agent B).[6] The resulting dynamics of disagreement exhibits countercyclical spikes, a pattern that does not arise when both agents use one of the two models with different but economically plausible parameters. When both agents use Agent A's model (e.g., Buraschi and Whelan (2013), Ehling et al. (2017), Buraschi, Trojani, and Vedolin (2014)), their assessment of uncertainty does not vary, which precludes spikes in disagreement and thus momentum. When both agents use Agent B's model (David (2008)), there exists combinations of parameters that give rise to counter-cyclical disagreement spikes, as in our model, but the parameter values estimated in David (2008) imply that disagreement never spikes and returns do not exhibit momentum. Intuitively, fitting the same model to the same data cannot

[6] Agent A's model is based on Detemple (1986, 1991), Brennan and Xia (2001), Scheinkman and Xiong (2003), and Dumas, Kurshev, and Uppal (2009), while Agent B's model is based on David (1997, 2008) and Veronesi (1999, 2000). The disagreement pattern varies depending on whether investors are overconfident (e.g., Scheinkman and Xiong (2003)), whether they have different initial priors (e.g., Detemple and Murthy (1994)), or whether they have dogmatic beliefs (e.g., Kogan et al. (2006)). See Xiong (2014) for a literature review.

produce sufficient heterogeneity in parameters across agents, and thus sufficiently different assessments of uncertainty, to generate momentum.

This paper is also related to the literature that studies the link between heterogeneous beliefs and momentum.[7] Banerjee, Kaniel, and Kremer (2009) show that heterogeneous beliefs generate price drift in the presence of higher-order differences in opinions. In contrast, we obtain momentum even when heterogeneous prior beliefs are commonly known. Ottaviani and Sorensen (2015) also obtain short-term momentum and long-term reversal through the combination of heterogeneous beliefs and wealth effects. Both in this paper and theirs, prices underreact because someone—Agent B in this paper and the marginal trader in theirs—reacts opposite to news. However, the mechanism in Ottaviani and Sorensen (2015) does not explain why this effect concentrates in bad times, which is the main focus of our paper.

The remainder of the paper is organized as follows. Section I presents and solves the model. Section II calibrates the model to the U.S. business cycle. Section III contains theoretical results on return predictability, while Section IV tests the predictions of the model. Section V concludes. Computational details are relegated to the Internet Appendix. (The Internet Appendix is available with the online version of the article on the Journal of Finance website.)

[7]Other theories of momentum include Berk, Green, and Naik (1999), Holden and Subrahmanyam (2002), Johnson (2002), Sagi and Seasholes (2007), Makarov and Rytchkov (2012), Vayanos and Woolley (2013), Biais, Bossaerts, and Spatt (2010), Cespa and Vives (2012), Albuquerque and Miao (2014), and Andrei and Cujean (2017).

I. The Model

We develop a dynamic general equilibrium in which investors use different models to estimate the business cycle. In this section, we describe the economy, we solve investors' learning and optimization problems, and we characterize the equilibrium stock price.

A. *The Economy and Models of the Business Cycle*

We consider an economy with an aggregate dividend that flows continuously over time. The market consists of two securities, a risky asset—the stock—in positive supply of one unit and a riskless asset—the bond—in zero net supply. The stock is a claim to the dividend process, δ, that satisfies

$$\mathrm{d}\delta_t = \delta_t f_t \mathrm{d}t + \delta_t \sigma_\delta \mathrm{d}W_t^\delta, \tag{1}$$

where W^δ is a standard Brownian motion under the objective probability measure, which governs the empirical realizations of dividends. The expected dividend growth rate f—henceforth the *fundamental*—is unobservable and follows the Ornstein-Uhlenbeck process

$$\mathrm{d}f_t = \kappa \left(\overline{f} - f_t \right) \mathrm{d}t + \sigma_f \mathrm{d}W_t^f, \tag{2}$$

where W^f is an independent standard Brownian motion under the objective probability measure. We assume that the system of equations in (1) and (2) represents the true data-generating process, an assumption we clarify in Section II.B.

The economy is populated by two agents, A and B, who consume the dividend and trade in the market. Agents understand that the fundamental affects the dividend they consume and the price of the assets they trade. Since the fundamental is unobservable, agents need to estimate it using the empirical realizations of dividends, the only source of information available. This information is meaningless, however, without a proper understanding of the data-generating process that governs dividends. Agents thus need to have a model in mind.

Agent A has the correct model in mind and understands that the system of processes in (1) and (2) generates the data she observes. In particular, she believes that the fundamental f evolves continuously over the business cycle according to (2), reverting to a long-term mean $\overline{f}$ at speed κ with constant uncertainty σ_f. In that respect, Agent A's beliefs define the beliefs of an outside observer—the econometrician—who understands the true data-generating process (e.g., Xiong and Yan (2010)).

Agent B believes that the fundamental follows a two-state continuous-time Markov chain and therefore uses the model

$$
\begin{aligned}
&\mathrm{d}\delta_t = f_t^B \delta_t dt + \sigma_\delta \delta_t \mathrm{d}W_t^B \\
&f_t^B \in \{f^h, f^l\} \text{ with generator matrix } \Lambda = \begin{pmatrix} -\lambda & \lambda \\ \psi & -\psi \end{pmatrix},
\end{aligned} \tag{3}
$$

where W^B is a Brownian motion under B's probability measure, which reflects her views about the data-generating process. Under Agent B's model, the fundamental f^B is either high f^h or low f^l. The economy moves from the high state to the low state with intensity $\lambda > 0$ and from the low state to the high state with intensity $\psi > 0$.

The financial economics literature has focused, largely, on the two types of model presented in equations (2) and (3) to forecast the growth rate of dividends. Agent A's model serves as a canonical model of dividend growth in the literature on heterogeneous beliefs (e.g., Scheinkman and Xiong (2003) and Dumas, Kurshev, and Uppal (2009)). The perspective of Agent B is based on the work of David (1997) and Veronesi (1999) and is closer to models used in the economics literature to forecast business cycle turning points.[8] Importantly, these models have always been considered separately. When considered jointly, these models lead to countercyclical spikes in disagreement among agents, the key feature of our framework.

B. *Bayesian Learning and Disagreement*

Agents learn about the fundamental by observing realizations of the dividend growth rate and, given the model they have in mind, they update their expectations accordingly. In doing so, they come up with an estimate of the fundamental, which we call the *filter*. We present the dynamics of the filters of Agents A and B in Proposition 1 below.

PROPOSITION 1:

1. *The filter of Agent A, $\widehat{f}_t^A = \mathbb{E}_t^{\mathbb{P}^A}\left[f_t^A\right]$, evolves according to*

$$d\widehat{f}_t^A = \kappa\left(\overline{f} - \widehat{f}_t^A\right)dt + \frac{\gamma}{\sigma_\delta}d\widehat{W}_t^A, \tag{4}$$

[8] See Hamilton (1994) and Milas, Rothman, and van Dijk (2006) for further details.

where $\gamma = \sqrt{\sigma_\delta^2\left(\sigma_\delta^2\kappa^2+\sigma_f^2\right)} - \kappa\sigma_\delta^2$ *denotes Agent A's steady-state posterior variance and* $\widehat{W}^A$ *is a Brownian motion under Agent A's probability measure* $\mathbb{P}^A$,

$$d\widehat{W}_t^A = 1/\sigma_\delta\left(d\delta_t/\delta_t - \widehat{f}_t^A dt\right). \tag{5}$$

2. *The filter of Agent B,* $\widehat{f}_t^B = \mathbb{E}_t^{\mathbb{P}^B}\left[f_t^B\right]$, *evolves according to*

$$d\widehat{f}_t^B = (\lambda+\psi)\left(f_\infty - \widehat{f}_t^B\right)dt + v(\widehat{f}_t^B)d\widehat{W}_t^B, \tag{6}$$

where $f_\infty = \lim_{t\to\infty}\mathbb{E}\left[f_t^B\right] = f^l + \frac{\psi}{\lambda+\psi}\left(f^h - f^l\right)$ *denotes the unconditional mean of the filter,* $v(\widehat{f}^B) = \frac{1}{\sigma_\delta}\left(\widehat{f}^B - f^l\right)\left(f^h - \widehat{f}^B\right)$ *defines Agent B's posterior variance, and* $\widehat{W}^B$ *is a Brownian motion under Agent B's probability measure* $\mathbb{P}^B$,

$$d\widehat{W}_t^B = 1/\sigma_\delta\left(d\delta_t/\delta_t - \widehat{f}_t^B dt\right). \tag{7}$$

Proof: See Internet Appendix Section I for Part *1* and Lipster and Shiryaev (2001) for Part *2*. ■

The two filters in equations (4) and (6) differ in terms of uncertainty. While Agent A's uncertainty γ/σ_δ does not fluctuate, Agent B's uncertainty $v(\cdot)$ does, reflecting her need to continuously reassess the likelihood of each state using dividends that themselves flow continuously. We analyze how this difference affects agents' disagreement in Section II.

Throughout the analysis we derive our results under the beliefs of Agent A—equivalently, those of the econometrician—which reflect the true data-generating process. We thus convert Agent B's views into those of Agent A through the following change of measure:

$$\left.\frac{\mathrm{d}\mathbb{P}^B}{\mathrm{d}\mathbb{P}^A}\right|_{\mathscr{F}_t} \equiv \eta_t = \exp\left[-\frac{1}{2}\int_0^t \frac{g_u^2}{\sigma_\delta^2}\mathrm{d}u - \int_0^t \frac{g_u}{\sigma_\delta}\mathrm{d}\widehat{W}_u^A\right], \tag{8}$$

where $g \equiv \widehat{f}^A - \widehat{f}^B$ represents agents' *disagreement* about the estimated fundamental. The change of measure in equation (8) implies that Agents A and B have equivalent perceptions of the world, a result we establish in Proposition 2.

PROPOSITION 2: *The probability measures $\mathbb{P}^A$ and $\mathbb{P}^B$ restricted to the filtration $\mathscr{F}_t = \sigma(\delta_u : u \le t)$ are equivalent for all $t \in \mathbb{R}_+$. In particular, at any time $t \in \mathbb{R}_+$ and for any threshold $c \in \mathbb{R}$, the cumulative distribution function of disagreement is bounded by*

$$\mathbb{P}^j\left(|g_t| \ge c\right) \le 2\sum_{i=1,2}\Phi\left(\frac{(-1)^i g_0 + \frac{\overline{m}}{\overline{\lambda}^j}\left(e^{\overline{\lambda}^j t} - 1\right) - c}{\sqrt{\frac{\overline{\sigma}^2}{2\overline{\lambda}^j}\left(e^{2\overline{\lambda}^j t} - 1\right)}}\right), \quad j = A, B \tag{9}$$

where Φ *denotes the normal cdf and the constants* $\overline{\lambda}^A$, $\overline{\lambda}^B$, $\overline{m}$, *and* $\overline{\sigma}$ *satisfy*

$$\lambda^j = \kappa + \mathbf{1}_{j=A}\left(\frac{f^h - f^l}{2\sigma_\delta}\right)^2 + \mathbf{1}_{j=B}\frac{\gamma}{\sigma_\delta^2}, \quad j = A, B$$

$$\overline{m} = \max\left\{|\kappa\overline{f} - (\lambda+\psi)f_\infty + (\lambda+\psi-\kappa)f^l|, |\kappa\overline{f} - (\lambda+\psi)f_\infty + (\lambda+\psi-\kappa)f^h|\right\}$$

$$\overline{\sigma} = \max\left\{\gamma, \left|\gamma - \frac{1}{4}\left(f^h - f^l\right)^2\right|\right\}.$$

It follows that the change of measure η *in* (8) *is a true martingale, which from Girsanov's theorem implies that the Brownian motions* $\widehat{W}^A$ *and* $\widehat{W}^B$ *satisfy the relation*

$$d\widehat{W}_t^B = d\widehat{W}_t^A + \frac{g_t}{\sigma_\delta}\,dt.$$

Proof: See Internet Appendix Section II. ■

Proposition 2 ensures that we can translate the perception of Agent B into that of Agent A. The economic relevance of this result is that, observing any finite history of data, Agent B cannot falsify her own model—if agents could sell a claim contingent on which model is correct, this claim would never pay off (except at an infinite horizon). This result relies critically on the dynamics of disagreement, which quantify the deviation between Agent B's beliefs and those implied by the true data-generating process. The bound in (9) ensures that, at any finite horizon, this deviation is never large enough for Agent B's model to become statistically inconsistent with the true data-generating process. That is, the distribution of disagreement never has a mass accumulating in the far tails. It follows that the process η remains strictly positive at any finite horizon (i.e., η is a true martingale). A direct implication of this result is that both agents survive at any finite time.

Over an infinite horizon, however, Agent B falsifies her model and vanishes asymptotically (see Dumas, Kurshev, and Uppal (2009) for a similar result). Since empiricists observe a century's worth of data at most, a relevant question is how long would it take for an empiricist to confidently reject Agent B's model. We elaborate on this question in Section II.B.

C. Equilibrium

Agents choose their portfolios and consumption plans to maximize their expected lifetime utility of consumption. They have power utility preferences defined by

$$U(c,t) \equiv e^{-\rho t}\frac{c^{1-\alpha}}{1-\alpha},$$

where $\alpha > 0$ is the coefficient of relative risk aversion and $\rho > 0$ the subjective discount rate.

Since markets are complete, we solve the consumption and portfolio choice problem of both agents using the standard martingale approach (Cox and Huang (1989)). We present the equilibrium risk-free rate, r^f, and the market price of risk perceived by Agent A, θ, in Proposition 3 below.

PROPOSITION 3: *The market price of risk under the viewpoint of Agent*

A, θ, and the risk-free rate, r^f, satisfy

$$\theta_t = \alpha\sigma_\delta + \frac{(1-\omega_t)}{\sigma_\delta} g_t \tag{10}$$

$$r_t^f = \rho + \alpha \widehat{f}_t^A - \frac{1}{2}\alpha(\alpha+1)\sigma_\delta^2 + (1-\omega_t)\, g_t \left(\frac{1}{2}\frac{\alpha-1}{\alpha\sigma_\delta^2}\omega_t g_t - \alpha \right), \tag{11}$$

where

$$\omega_t = \frac{(1/\phi_A)^{1/\alpha}}{(1/\phi_A)^{1/\alpha} + (\eta_t/\phi_B)^{1/\alpha}} \tag{12}$$

denotes the consumption share of Agent A and where ϕ_A and ϕ_B are the Lagrange multipliers associated with the budget constraints of Agents A and B, respectively.

Proof: See Internet Appendix Section III. ■

We focus our discussion on the market price of risk in equation (10), which is a key determinant of return predictability. The risk-free rate in equation (11) is discussed in detail in David (2008) and Buraschi and Whelan (2013). The market price of risk is the product of the diffusion of Agent A's consumption growth,

$$\sqrt{\frac{1}{\mathrm{d}t}\mathrm{var}_t^{\mathbb{P}^A}\left(\frac{\mathrm{d}c_{At}}{c_{At}}\right)} = \sigma_\delta + \frac{1-\omega_t}{\alpha\sigma_\delta} g_t, \tag{13}$$

and her coefficient of risk aversion, α. Agent A wants to be compensated for holding assets that co-vary positively with her consumption growth and is willing to pay a premium for holding assets that co-vary negatively with her consumption growth. In particular, her consumption varies either when dividends fluctuate or when her consumption share fluctuates. Hence, the market price of risk is proportional to the diffusion of dividend growth and the diffusion of the growth of her consumption share. The market price of risk then scales with risk aversion, reflecting that the more risk-averse Agent A is, the more she wants to hedge.

The consumption share of Agent B, $1 - \omega$, determines the extent to which disagreement affects the market price of risk. Suppose disagreement is positive today (Agent B is pessimistic). A positive dividend shock tomorrow then indicates that Agent B's beliefs were inaccurate—the likelihood of Agent B's model relative to the true data-generating process, η, decreases and the consumption share of Agent A increases through (12). In this case, Agent A wants to be rewarded for holding assets that co-vary positively with the dividend shock and thus with her consumption share. Moreover, as the consumption share of Agent B decreases, the risk of Agent A's consumption share growth decreases: were Agent A the only agent populating the economy, disagreement would become irrelevant and Agent A would be compensated for holding assets that are positively correlated with dividend growth only. This mechanism applies symmetrically when disagreement is negative.

In Proposition 4 we conclude the equilibrium description by providing the equilibrium stock price S, the derivation of which follows the methodology in Dumas, Kurshev, and Uppal (2009).

PROPOSITION 4: *Assuming that the coefficient of relative risk aversion* α

is an integer, there exists an equilibrium in which the stock price satisfies

$$\frac{S_t}{\delta_t} = \omega_t^{\alpha} \frac{S_t}{\delta_t}\bigg|_{O.U.} + (1-\omega_t)^{\alpha} \frac{S_t}{\delta_t}\bigg|_{M.C.} + \omega_t^{\alpha} \sum_{j=1}^{\alpha-1} \left(\begin{array}{c} \alpha \\ j \end{array} \right) \left(\frac{1-\omega_t}{\omega_t} \right)^j F^j \left(\widehat{f}_t^A, g_t \right), \tag{14}$$

where $\frac{S}{\delta}\big|_{O.U.}$ *and* $\frac{S}{\delta}\big|_{M.C.}$ *denote the prices that prevail in a representative-agent economy populated by Agents A and B, respectively. The functions* $F^j \left(\widehat{f}^A, g \right)$ *represent price adjustment for disagreement. Under the calibration of Section II.B, this equilibrium is unique.*

Proof: See Proposition 3 in David (2008) for existence and Internet Appendix Section IV. ■

The price in equation (14) has three terms. When agents have logarithmic utilities ($\alpha = 1$), only the first two terms are relevant and the price is just an average of the prices that obtain in representative-agent economies populated by Agents A and B. When agents are more risk-averse than a logarithmic agent ($\alpha > 1$), the third term becomes relevant. It describes the joint effect of fundamental and disagreement on the price.

II. Dynamics of Disagreement and Calibration

In this section, we analyze the dynamics of disagreement that arise when one agent uses a discrete-state model and the other a continuous-state model. In Section II.A, we provide restrictions on model parameters under which this difference in models creates countercyclical spikes in disagreement, the key feature driving our results. We then calibrate both models and let the data decide whether these parameter restrictions are satisfied (Section II.B). We illustrate the resulting pattern of disagreement in Section II.C: agents' expectations polarize in bad times, adjust at different speeds in normal times, and move together in good times, leading to countercyclical spikes in disagreement.

A. A Mechanism for Countercyclical Spikes in Disagreement

In our framework, agents use different models and thus learn differently, as apparent from equations (4) and (6). First, Agent B's discrete-state model implies state-dependent uncertainty $v(\widehat{f}^B)$, whereas Agent A's continuous-state model implies constant uncertainty γ/σ_δ. We show that this difference in agents' assessment of uncertainty causes Agent B's views to become self-reinforcing (centrifugal) under the true probability measure $\mathbb{P}^A$, a mechanism that can lead disagreement among agents to spike. Second, agents' views evolve at different speeds ($\lambda+\psi$ and κ) and have different long-term means (f_∞ and $\overline{f}$). This difference in persistence and long-term mean has an asymmetric effect on disagreement, which can lead disagreement spikes to concentrate in bad states of the economy.

To guide the analysis, we introduce the concepts of "adjustment speed" and "polarization of beliefs," which we provide in Definition 1.

DEFINITION 1: *The speed at which Agent i updates her expectations conditional on the initial value of her filter, Σ_t^i, is the rate of change of her* *average filter,* $\mathbb{E}_0^{\mathbb{P}^A}[\widehat{f}_t^i | \widehat{f}_0^i = x_0]$*, over time*

$$\Sigma_t^i := \frac{d}{dt}\mathbb{E}_0^{\mathbb{P}^A}[\widehat{f}_t^i | \widehat{f}_0^i = x_0], \quad i = A, B.$$

Furthermore, a polarization of beliefs occurs at time t *when* $sign(\Sigma_t^A) \neq sign(\Sigma_t^B)$*.*

The change of measure of Proposition 2 modifies the speed at which Agent B learns. Expressing Agent B's expectations under the true measure $\mathbb{P}^A$ introduces a nonlinear term in Agent B's adjustment speed through the state-dependent uncertainty $v(\widehat{f}^B)$ of her filter:

$$\Sigma_t^B = \mathbb{E}_0^{\mathbb{P}^A}\left[(\lambda + \psi)\left(f_\infty - \widehat{f}_t^B\right) + \underbrace{\frac{1}{\sigma_\delta}\left(\widehat{f}_t^A - \widehat{f}_t^B\right) v(\widehat{f}_t^B)}_{\text{change of measure}}\right]. \tag{15}$$

To understand how the change of measure affects the speed at which Agent B learns, we make two simplifying assumptions, which are consistent with our calibration in Section II.B.

ASSUMPTION 1: *Agent B's model has symmetric states* $f^h = -f^l \equiv f^\star$ *around zero.*

ASSUMPTION 2: *The volatility of dividends* σ_δ *is such that terms of order* $o(\sigma_\delta^4)$ *are negligible.*

We next then perform a second-order approximation of the change of measure in (15) and denote by $\widetilde{\Sigma}_t^B$ the resulting adjustment speed for Agent B, which we present in Proposition 5.

PROPOSITION 5: *Under Assumptions 1 and 2, the second-order approximation of Agent B's adjustment speed in a neighborhood of* $t \approx 0$ *satisfies*

$$\widetilde{\Sigma}_t^B = \underbrace{(1-\pi(x_0,t))\overbrace{(\lambda+\psi)(f_\infty - x_0)}^{\equiv\, \Sigma_t^B \text{ under } \mathbb{P}^B} + \pi(x_0,t)\overbrace{\kappa(\overline{f}-x_0)}^{\equiv \Sigma_t^A}}_{\text{weighted adjustment speed}} + \underbrace{2\frac{\pi(x_0,t)^2}{t}x_0}_{\text{centrifugal effect } \equiv C(x_0)}, \tag{16}$$

where $\pi(x_0,t) = v(x_0)t/\sigma_\delta$.

Proof: See Internet Appendix Section V. ■

The nonlinear change of measure in (15) has two effects on Agent B's adjustment speed. First, it adjusts the dynamics of Agent B's expectations to make them consistent with the correct expectations—those of Agent A. Specifically, the first term in (16) is a weighted average of agents' adjustment speeds under their respective probability. Agent B's uncertainty drives the weight $\pi(x_0,t)$, which reaches its maximum, $\pi(f_m,t) = (f^\star/\sigma_\delta)^2 t$, when Agent B's filter is equal to

$$f_m = \frac{1}{2}f^l + \frac{1}{2}f^h (\equiv 0 \text{ under Assumption 1}),$$

the point at which Agent B is most uncertain about the state of the economy. For ease of interpretation, assume that $t = (\sigma_\delta/f^\star)^2$ so that $\pi(x_0,(\sigma_\delta/f^\star)^2) \in [0,1]$ in all states. It follows that the first term in equation (16) causes agents' disagreement to increase when Agent B is confident that the economy is in a bad or a good state.

Second, the change of measure causes Agent B's expectations to become

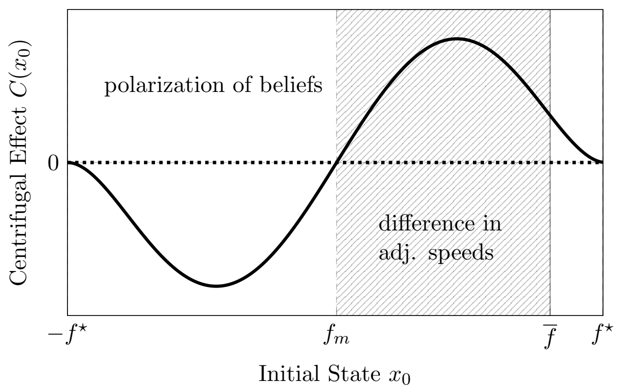

Figure 1. Centrifugal effect on the speed at which Agent B learns. This figure illustrates the centrifugal effect on Agent B's adjustment speed (see Proposition 5) as a function of the initial value of Agent B's filter.

self-reinforcing (centrifugal) outward, f_m, through the second term, $C(x_0)$, in equation (16). To illustrate this effect, in Figure 1 we plot this term as a function of the initial value of Agent B's filter. As Agent B's expectations rise above f_m, they increase at an accelerated rate, which generates a *difference in adjustment speeds* across the two agents in good states. As Agent B's expectations drop below f_m, they decrease at a faster rate, which creates a *polarization of beliefs* across the two agents in bad states. In other words, good news in good states causes Agent B to revise her expectations upwards sharply. In contrast, good news in bad states always corroborates Agent A's view that the economy is recovering, but is never "good enough" to convince Agent B that economic conditions are improving.

By amplifying disagreement in good and bad states, the centrifugal effect can produce *spikes* in short-term disagreement. Specifically, the maximum magnitude of this effect is

$$\max_{x_0 \in [-f^\star, f^\star]} |C(x_0)| = \frac{8}{5^{3/2}} \frac{f^\star}{\sigma_\delta} v(f^\star/\sqrt{5}), \text{ around time } t \approx \left(\frac{\sigma_\delta}{f^\star}\right)^2.$$

It follows that the larger the ratio of the distance between Agent B's states to the volatility of dividends, the larger and sharper the changes in disagreement. For instance, a ratio of $f^\star/\sigma_\delta \equiv 10$ implies that disagreement can change Agent B's uncertainty, v, by up to sevenfold in a matter of days ($t \approx 0.01$). The economic relevance of such sudden and large changes in disagreement is that they can generate stock return momentum in our model, a key result that we present in Section III. In that respect, we reserve the terminology "disagreement spikes" to refer to changes in disagreement that are sufficiently large and sharp to produce momentum. Since the ratio $f^\star/\sigma_\delta$ determines the magnitude of changes in disagreement, the existence of disagreement spikes in the model places an implicit restriction on this ratio.

CONDITION 1: *Under Assumptions 1 and 2, a necessary condition for disagreement spikes to arise in the model is that the ratio $f^\star/\sigma_\delta$ be sufficiently large.*

Importantly, without state-dependent uncertainty, the centrifugal effect in equation (16) is absent. Consider, for instance, the specification under which both agents use Agent A's linear continuous-state model (e.g., Buraschi and Whelan (2013), Ehling et al. (2017), and Buraschi, Trojani, and Vedolin (2014)). In this case the change of measure in equation (15) only has a linear effect on Agent B's adjustment speed, making it challenging to generate sharp increases in disagreement. A discrete-state model instead implies state-dependent uncertainty, which could lead disagreement to spike.

Since the centrifugal effect acts symmetrically on the magnitude of disagreement spikes across good and bad states, any asymmetry—procyclicality or countercyclicality—must arise through the weighted average in equation (16). Notice first that the speed at which disagreement increases is the difference between agents' adjustment speeds, which satisfies

$$\Sigma_t^A - \widetilde{\Sigma}_t^B = (1 - \pi(x_0, t)) \left(\underbrace{(\lambda - \psi) f^\star + \kappa \overline{f}}_{\text{asymmetric effect}} - (\kappa - \psi - \lambda) x_0 \right) - C(x_0). \tag{17}$$

Given that the centrifugal effect makes disagreement spikes positive in bad states and negative in good states, a necessary condition for disagreement spikes to be stronger in bad states is that the asymmetric effect in equation (17) be positive. When this effect is positive, it accentuates the polarization of beliefs in bad states and dampens the difference in adjustment speeds in good states. This condition for countercyclical spikes imposes a restriction on the persistence of good states relative to bad states, which we provide in Condition 2.

CONDITION 2: *Under Assumptions 1 and 2, a necessary condition for disagreement spikes to be countercyclical is that* $\psi - \lambda < \kappa \frac{\overline{f}}{f^\star}$.

Conditions 1 and 2 jointly define parameter restrictions under which disagreement exhibits countercyclical spikes in our framework. The mechanism for disagreement spikes that we highlight is not unique to our framework—a similar mechanism applies when both agents use Agent B's model with different parameters (David (2008)). Under both model specifications, there exists combinations of parameters that give rise to disagreement spikes and countercyclicality. However, these effects must arise through an economically plausible mechanism, not through an arbitrary choice of parameters. We therefore let the data decide whether the parameter restrictions of Conditions 1 and 2 are satisfied and whether these estimated effects distinctly identify our specification, a matter we investigate next.

B. *Calibration and Model Fit to the U.S. Economy*

In our model, agents use a single source of information—the time series of dividends—to update their expectations. As a proxy for the dividend stream, we use the S&P 500 dividend time series recorded at a monthly frequency from January 1871 to November 2013, which we obtain from Robert Shiller's website. Looking back to the 19th century allows us to cover a large number of business cycle turning points, but obviously adds strong seasonality effects (Bollerslev and Hodrick (1992)). To reduce these effects, we apply the filter developed by Hodrick and Prescott (1997) to the time series of dividends.

We assume that both agents observe the S&P 500 dividend time series after seasonalities have been smoothed out. Since these data are available monthly, agents need to first estimate a discretized version of their model

Table I
Parameter Calibration

This table reports the estimated parameters of the continuous-time model of Agent A and Agent B. Standard errors (computed using the delta method) are reported in brackets and statistical significance at the 10%, 5%, and 1% levels is denoted by * , **, and ***, respectively. The last three rows report our choice of preference parameters and initial consumption shares ($\phi_A = \phi_B \Leftrightarrow \omega_0 = 1 - \omega_0 = 0.5$).

Parameter	**Symbol**	**Value**
Volatility of Dividend Growth	σ_δ	0.0225*** (3.94×10^{-4})
Mean-Reversion Speed of f^A	κ	0.1911*** (0.0264)
Long-Term Mean of f^A	$\bar{f}$	0.0630*** (0.0083)
Volatility of f^A	σ_f	0.0056*** (2.34×10^{-4})
High State of f^B	f^h	0.0794*** (0.0032)
Low State of f^B	f^l	−0.0711*** (0.0038)
Intensity of f^B: High to Low	λ	0.3022 (1.1294)
Intensity of f^B: Low to High	ψ	0.3951 (1.1689)
Relative Risk Aversion	α	2
Subjective Discount Rate	ρ	0.01
Lagrange Multipliers	$\phi_A = \phi_B$	1

by Maximum Likelihood.[9] Agents then map the parameters they estimated into their continuous-time model. In doing so, agents obtain the parameter values presented in Table I. All discussions and results that follow are based on these parameter values. For convenience, we discuss the methodological details in Internet Appendix Section VI.

The parameters of Table I show that both models produce distinct in-

[9]See Hamilton (1994) for the likelihood function of each model. We present the estimated parameters, their standard errors, and their statistical significance in Table IA.I in Internet Appendix Section VI.

terpretations of the data. First, Agent A finds a low reversion speed κ and therefore concludes that the fundamental is persistent. Second, Agent B finds that the high state f^h and the low state f^l are symmetric around zero, consistent with Assumption 1, but that transition intensities between the two states are asymmetric—expansions are more persistent than recessions ($\psi > \lambda$). The values of the transition intensities ψ and λ further imply that Agent B's filter reverts about 3.5 times faster than that of Agent A, and that the long-term mean of Agent B's filter ($f_\infty \approx 0.014$) is significantly lower than that of Agent A ($\overline{f} = 0.063$).

This calibration allows us to determine which model better fits historical data by applying a model selection method such as the Akaike Information Criterion (AIC).[10] Because an Ornstein-Uhlenbeck process is more versatile than a two-state Markov chain, the information criterion favors Agent A's model.[11] This fact ultimately justifies our assumption that Agent A's model in equations (1) and (2) is the true data-generating process and thus our choice of computing and analyzing the equilibrium under Agent A's probability measure $\mathbb{P}^A$. However, it does not necessarily question Agent B's rationality, as we now illustrate.

To evaluate whether Agent B's model is plausible when the data are actually generated by Agent A's model, we conduct the following experiment. We simulate a discrete time series $\{y_\tau\}_{\tau=1}^n$ of n monthly dividend growth using the true data-generating process in equations (1) and (2) and

[10]The Akaike Information Criterion is defined as $AIC = 2K - 2\log(L)$, where K is the number of parameters estimated and L the likelihood function. The smaller the criterion, the better the model.

[11]The AIC for Agent A and B's models are $AIC_{\text{Agent A}} = -1.7229 \times 10^4$ and $AIC_{\text{Agent B}} = -1.2154 \times 10^4$, respectively.

the parameter values of Table I. We then consider an econometrician whose goal is to determine which model—Agent A's or Agent B's—is most likely to have generated this time series. Using the time series of dividend growth, the econometrician first estimates by Maximum Likelihood the parameters $\Theta^A \equiv (\sigma_\delta, \kappa, \bar{f}, \sigma_f)$ of Agent A's model in (4) and (5) and the parameters $\Theta^B \equiv (\sigma_\delta, \lambda, \psi, f^h, f^l)$ of Agent B's model in (6) and (7). Following Vuong (1989) she then constructs a likelihood-ratio test of the two (nonnested) models using the statistic LR:

$$LR \equiv n^{-1/2} \frac{\sum_{i=1}^{n} \log \frac{h^A(y_{t+1}|\widehat{f}_t^A;\Theta^A)}{h^B(y_{t+1}|\widehat{f}_t^B;\Theta^B)}}{\sqrt{\frac{1}{n}\sum_{i=1}^{n}\left[\log \frac{h^A(y_{t+1}|\widehat{f}_t^A;\Theta^A)}{h^B(y_{t+1}|\widehat{f}_t^B;\Theta^B)}\right]^2 - \left[\frac{1}{n}\sum_{i=1}^{n}\log \frac{h^A(y_{t+1}|\widehat{f}_t^A;\Theta^A)}{h^B(y_{t+1}|\widehat{f}_t^B;\Theta^B)}\right]^2}}$$

$$\xrightarrow{d} \mathscr{N}(0,1),$$

where $\mathscr{N}$ denotes the normal distribution and h^A and h^B denote the density associated with Agent A's and B's model, respectively. If $LR > \Phi^{-1}(p)$, the econometrician rejects the null that both models are equivalent in favor of Agent A's model at the $p\%$ confidence level.

Figure 2 plots the rejection frequency (at the 90%, 95%, and 99% confidence level) over time, based on 1,000 simulated time series of monthly dividends, each of length 1,000 years. Figure 2 shows that the econometrician would need to observe more than three centuries' worth of data to correctly reject the null in favor of Agent A's model more than 50% percent of time, even at the 90% confidence level. After 1,000 years, the econometrician would still fail to reject the null more than 10% of the time. Considering that empiricists observe a century's worth of data, we conclude from this experiment that Agent B's model is unlikely to be confidently rejected over a time period that is empirically relevant.

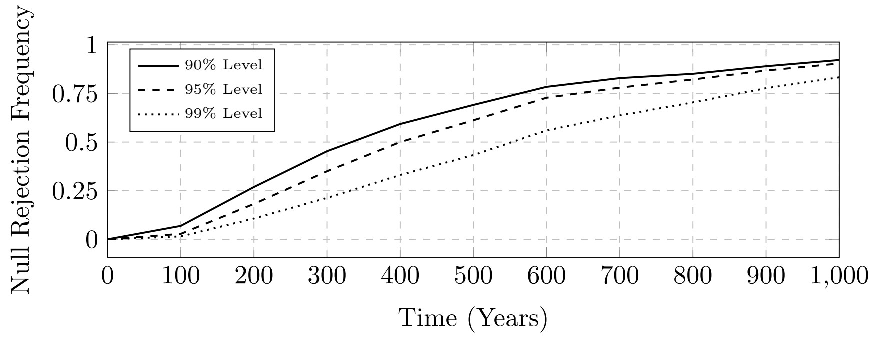

Figure 2. Rejection frequency of Agent B's model. This figure plots the rejection frequency of Agent B's model over time. Each line corresponds to the rejection frequency of the null in favor of Agent A's model at the 90%, 95%, and 99% confidence level, respectively. The rejection frequency is based on 1,000 simulations of monthly dividend growth rates observed over 1,000 years.

That Agent B's model is not rejected quickly directly implies that Agent B will not disappear quickly. To show this, in Figure 3 we plot representative paths of the relative likelihood of Agent B's model, η, and the corresponding consumption share of Agent A, ω. Illustrating the result of Proposition 2 that the likelihood η is a true martingale, Figure 3 shows that η cycles, thus allowing Agent B to remain influential in the model. Although Agent A's consumption grows on average (given that she uses the correct model), Figure 3 shows that Agent B sometimes enjoys a larger share of consumption. This phenomenon occurs when, by luck, Agent B's model produces a succession of forecasting errors that are smaller than those of Agent A. For instance, when both agents underestimate the fundamental but Agent B

does so less, the likelihood of Agent B's model and hence her consumption share increase. If this phenomenon occurs sufficiently many times in a row, Agent B's consumption may dominate Agent A's, as Figure 3 shows.

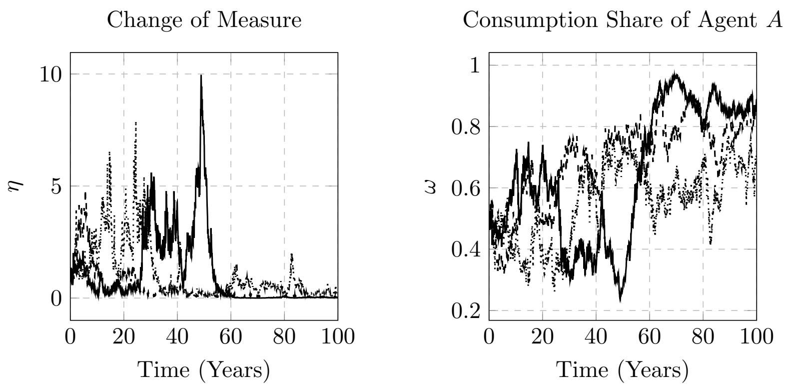

Figure 3. Representative paths of the likelihood of Agent B's model and Agent A's consumption share. This figure plots in separate panels representative simulated paths of the likelihood of Agent B's model, η, and the corresponding consumption share of Agent A, ω. Each simulation is based on the calibration of Table I.

Finally, another important aspect of the calibration in Table I is that it satisfies Conditions 1 and 2, thus allowing for countercyclical spikes in disagreement. Notice first that the low volatility of dividends in Table I is consistent with Assumption 2 and implies that the ratio of Condition 1 can produce disagreement that changes Agent B's uncertainty by up to fivefold within days. In Section III we show that these disagreement spikes are sufficiently large to produce momentum in stock returns (Condition 1). To emphasize the particularity of this result, consider the specification in which both agents use Agent B's model with different parameters (David (2008)). The parameter values estimated in David (2008) then imply that disagreement never spikes and stock returns do not exhibit momentum under this specification. In other words, fitting the same discrete-state model to the data does not produce sufficient heterogeneity in parameters to generate momentum in the model, hence the relevance of our specification. Second, the relative persistence of good and bad states meets Condition 2. We now show that this condition is sufficient to make disagreement spikes countercyclical under our calibration.

C. Estimated Dynamics of Disagreement

In this section we illustrate how the mechanism of Section II.A and the calibration of Section II.B combine to produce countercyclical spikes in disagreement, consistent with observed patterns among forecasters (e.g., Kandel and Pearson (1995) and Patton and Timmermann (2010)). We first define three regimes of the economy, which we use to describe the different phases of the business cycle throughout the analysis. Based on the results of Section II.A, the relevant business cycle turning point for Agent A is her long-term mean, $\bar{f}$, while that for Agent B is the point of maximum uncertainty f_m. Accordingly, we say that the economy is going through *good times* when agents' expectations are above $\bar{f}$, while the economy is going through *bad times* when agents' expectations are below f_m. Otherwise, when expectations lie between f_m and $\bar{f}$, we say the economy is in *normal times*. To emphasize that future disagreement—as opposed to current disagreement—drives our result, we set current disagreement to zero in each case. As a convention, we assume that agents' filters start in the middle of

Table II

Definition of Regimes

This table describes the three different regimes of the economy—good, normal, and bad times.

Good Times	**Normal Times**	**Bad Times**
$(\widehat{f}^A, \widehat{f}^B) \in [\bar{f}, f^h], \;\; g = 0$	$(\widehat{f}^A, \widehat{f}^B) \in [f_m, \bar{f}], \;\; g = 0$	$(\widehat{f}^A, \widehat{f}^B) \in [f^l, f_m], \;\; g = 0$

each interval.[12] Table II summarizes the definition of the three regimes.

To illustrate the behavior of disagreement over the business cycle, in Figure 4 we plot Agent A and B's average filter over time, with each regime highlighted in a separate panel. In good times (the left panel), both agents adjust their views at comparable speeds and disagreement exhibits little variation. In normal times (the middle panel), the difference in adjustment speed between Agents A and B becomes apparent. Although both agents expect economic conditions to improve—both filters move upwards on average—Agent B adjusts her expectations sharply due to the centrifugal effect in equation (16). This difference in adjustment speeds causes disagreement to spike in the short term.

In bad times (the right panel), beliefs polarize in the short term through the centrifugal effect in equation (16). Good news in bad times always corroborates Agent A's optimistic views, but is never good enough to invalidate Agent B's pessimistic views, reinforcing her beliefs that economic conditions are deteriorating. This polarization of opinions creates disagreement spikes

[12]The dynamics of disagreement are consistent within each region, irrespective of the starting point within each region, except in two knife-edge intervals, $(0.063, 0.067)$ and $(-0.056, -0.0711)$, around $\bar{f}$ and f^l, respectively. In Internet Appendix Section XI.A, we explain why disagreement dynamics differ, but show that our results remain unaffected, within these intervals.

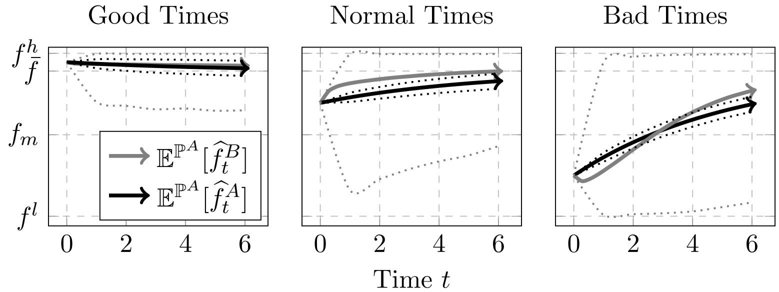

Figure 4. Filtered dynamics in three states of the economy. The solid black and solid gray lines represent Agent A and B's average filters ($\mathbb{E}^{\mathbb{P}^A}[\widehat{f}_t^A]$ and $\mathbb{E}^{\mathbb{P}^A}[\widehat{f}_t^B]$), respectively. The dotted lines depict the corresponding 90% confidence intervals. Filters are plotted against time. Each panel corresponds to a specific state of the economy: good, normal, and bad times.

that are strongest in bad times through the asymmetric effect in equation (17). This asymmetry arises because the high state of Agent B's model is more persistent than its low state. As a result, Agent B's uncertainty increases as economic conditions deteriorate (see the confidence interval in each panel of Figure 4), making the effect of the change of measure in equation (15) stronger in bad times. In the long run, Agent B eventually realizes that the economy is recovering and rapidly catches up with Agent A, thus becoming overly optimistic. Disagreement therefore evaporates in the medium term and regenerates in the long term through the difference in adjustment speeds.

III. Disagreement Driving Stock Return Predictability

In this section we show that countercyclical spikes in disagreement cause stock return predictability to concentrate in bad times. In our model fluctuations in disagreement command a risk premium: contemporaneous excess returns decrease with the *square* of disagreement (Section III.A). The sign of disagreement determines only the direction in which excess returns adjust in the future (Section III.B). In bad times, Agent B reacts opposite to news, beliefs polarize, and future excess returns underreact. In normal times, Agent B overreacts to news, beliefs adjust at different speeds, and future excess returns overreact. In good times, beliefs move together, disagreement is nearly constant, and future excess returns adjust immediately. Since spikes in disagreement—positive *or* negative—persistently raise the risk premium, both under- and overreaction create positive serial correlation in excess returns in the short term, a phenomenon known as time-series momentum (Section III.C). That disagreement spikes are countercyclical makes this phenomenon stronger in bad times.

To understand how disagreement, contemporaneous excess returns, and future excess returns are related, it is useful to describe trading strategies. Following Dumas, Kurshev, and Uppal (2009), Proposition 6 decomposes Agent A's strategy into two components.

PROPOSITION 6: *The number of shares, Q, that Agent A holds can be decomposed into a myopic portfolio, M, and a hedging portfolio, H, according* *to*

$$Q_t = M_t + H_t = \frac{V_t}{\alpha \sigma_t S_t} \theta_t + \frac{\alpha - 1}{\alpha \sigma_t S_t} E_t^{\mathbb{P}^A} \left[\int_t^\infty \frac{\xi_s}{\xi_t} c_{As} \left(\frac{\mathscr{D}_t \xi_s}{\xi_s} - \frac{\mathscr{D}_t \xi_t}{\xi_t} \right) ds \right], \tag{18}$$

where ξ denotes Agent A's state-price density, which is given by

$$\xi_t = e^{-\rho t} \delta_t^{-\alpha} \left[(1/\phi_A)^{1/\alpha} + (\eta_t/\phi_B)^{1/\alpha} \right]^\alpha, \tag{19}$$

and V denotes Agent A's wealth, which we provide in Internet Appendix Section VII.

Proof: See Internet Appendix Section VII. ■

Agent A's portfolio in (18) tells us how she trades on return predictability. While the first part, M, is a myopic demand through which Agent A seeks to extract the immediate Sharpe ratio, θ, the second term, H, is a hedging demand through which she exploits return predictability. To see this, notice that the hedging demand in (18) incorporates Agent A's outlook on future returns through the response of future state-price densities to a shock occurring today, $\mathscr{D}_t \xi_s$. This response represents Agent A's attempt to predict future returns. We now analyze how disagreement affects both contemporaneous and future excess returns and then derive implications for the serial correlation of excess returns.

A. Disagreement and Contemporaneous Excess Returns

Our goal is to determine the relation between contemporaneous excess returns and the state variables of the model. A key result of this section is that contemporaneous excess returns decrease with the square of disagreement. Contemporaneous (expected) excess returns, $\mu - r^f = \sigma\theta$, are the product of the market price of risk, θ, which we discuss in Section I.C, and the diffusion of stock returns, σ:

$$\sigma_t = \overbrace{\sigma_\delta + \frac{\gamma}{\sigma_\delta}\frac{1}{S_t}\left(\underbrace{\frac{\partial S}{\partial \widehat{f}^A}}_{<0} + \underbrace{\frac{\partial S}{\partial g}}_{>0}\right)}^{\text{component 1: } A\text{'s uncertainty}} \overbrace{-(1-\omega_t)v\left(\widehat{f}_t^A - g_t\right)\frac{1}{S_t}\underbrace{\frac{\partial S}{\partial \widehat{g}}}_{>0}}^{\text{component 2: } B\text{'s uncertainty}}$$
$$+ \overbrace{\frac{g_t(1-\omega_t)\omega_t}{\alpha\sigma_\delta^2}\frac{1}{S_t}\underbrace{\frac{\partial S}{\partial \omega}}_{<0}}^{\text{component 3: } A\text{'s cons. share}}, \tag{20}$$

where $\widehat{g} \equiv (1-\omega)g$ denotes the "consumption-weighted" disagreement (see Internet Appendix Section VIII).

We first explain the sign of the price sensitivities in equation (20). The price decreases with the fundamental through the income effect (when risk aversion is larger than one (Veronesi (2000)): anticipating that an increase in the fundamental today leads to higher consumption tomorrow, Agent A decreases her savings, which decreases the price. In contrast, the price increases with disagreement through the substitution effect: an increase in disagreement implies that Agent A is optimistic and therefore increases her stock holdings, which leads to a price increase. Finally, an increase in Agent A's consumption share implies a decrease in the risk of Agent A's consumption growth through (13), to which she responds by decreasing her hedging demand, leading to a price decrease.

We now determine how the diffusion of stock returns depends on the state variables of the model. The diffusion in equation (20) has three components, which we plot in separate panels in Figure 5. Because Agent A's uncertainty is constant, the first diffusion component in the left panels of Figure 5 is nearly constant. In contrast, Agent B's state-dependent uncertainty causes the second component (the middle panels) to decrease with disagreement. This decreasing relation steepens when Agent A's filter increases or when her consumption share decreases. This interaction between Agent A's views and Agent B's uncertainty raises stock return volatility. Furthermore, fluctuations in Agent A's consumption share induce a decreasing relation between the third component and disagreement (the right-hand panels), which steepens as agents' consumption shares equalize. Overall, our discussion implies that the diffusion of stock returns can be approximated as

$$\sigma_t \approx A_0 + A_1 g_t \widehat{f}_t^A (1 - \omega_t) + A_2 g_t \omega_t (1 - \omega_t), \tag{21}$$

where the signs of the constants $A_0 \approx 0$, $A_1 < 0$, and $A_2 < 0$ follow from price sensitivities.

Using Assumption 2 that σ_δ is small, we finally obtain a relation for excess returns by multiplying the diffusion in equation (21) with the market

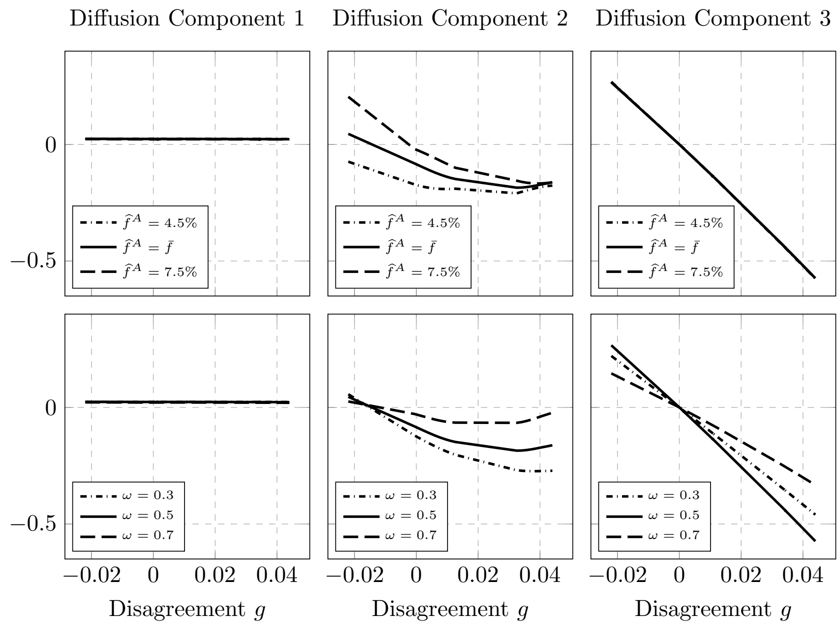

Figure 5. Decomposition of the stock return diffusion. The upper and lower panels show the three components of stock return diffusion against disagreement for different values of the fundamental and Agent A's consumption share, respectively. If not stated otherwise, we set $\delta = 1$, $\widehat{f}^A = \bar{f}$, and $\omega = 0.5$. The chosen range for disagreement corresponds to its 90% confidence interval.

price of risk in equation (10):

$$\mu_t - r_t^f \approx \text{constant}(\approx 0) + \frac{A_1}{\sigma_\delta}\widehat{f}_t^A \left((1-\omega_t)g_t\right)^2 + \frac{A_2}{\sigma_\delta}\omega_t \left((1-\omega_t)g_t\right)^2 . \tag{22}$$

Equation (22) accurately approximates excess returns and serves as the basis for our empirical specifications in Section IV.[13] This relation has two main implications for excess returns. First, contemporaneous excess returns

[13]Simulations show that regressing excess returns, $\mu - r^f$, on both $\widehat{f}^A \left((1-\omega)g\right)^2$ and $\omega \left((1-\omega)g\right)^2$ yields a R^2 of 95%, while regressing it on $\omega \left((1-\omega)g\right)^2$ yields a R^2 of 85%. See Internet Appendix Section VIII for further details.

are hump-shaped in disagreement and decreasing in the square of disagreement. Second, as the right-hand panels of Figure 5 show, the main source of fluctuations in excess returns is the last term in equation (22).

B. Disagreement, Future Excess Returns, and Reaction to News

We now show that future excess returns react to contemporaneous news in a way that is predictable if disagreement spikes in the short term. In bad times, news polarizes opinions, causing returns to persistently underreact. In normal times, news exacerbates the difference in beliefs' adjustment speed, causing returns to persistently overreact. In good times, news provokes an immediate return adjustment. Since disagreement spikes are countercyclical, news content (Tetlock (2007)) better predicts future returns in bad times (Garcia (2013)).

Agent A's hedging demand in (18) shows that forecasting future returns in our model involves computing the response of the future state-price density to a shock occurring today. Hence, as emphasized in Dumas, Kurshev, and Uppal (2009), the concept of future returns that is relevant for portfolio choice is the future response, $\mathscr{D}_t \xi_s / \xi_s$, of the state-price density, as opposed to the usual multiperiod rate of return, ξ_t / ξ_s. We accordingly refer to "future returns" and "the response of the state-price density" indifferently and denote future returns from time t to horizon $s > t$ by $R(t, s) = \mathscr{D}_t \xi_s / \xi_s$. This concept can be mapped into the empirical measures of Tetlock (2007) and Garcia (2013): the average, $\mathbb{E}_t^{\mathbb{P}^A}[R(t, s)]$, is the regression coefficient of news arriving at time t on future stock returns from time t to s.

We start with an illustration of the concept of future response. To compute a response to a news shock, Agent A first considers a trajectory of the Brownian innovation $\widehat{W}^A$, the only source of news in our model. She then contemplates a news surprise today (an initial perturbation of the news trajectory), while keeping the news trajectory otherwise unchanged, as illustrated in the upper panel of Figure 6. Each news trajectory, both perturbed and unperturbed, is associated with a trajectory of the state-price density, as illustrated in the lower panel. The difference, $\mathscr{D}\xi$, between these trajectories (the shaded area) captures the reaction of the state-price density to the news surprise relative to the state-price density that would have prevailed had there been no news surprise. Because Agent A is interested in all possible trajectories, she computes an average reaction $E[\mathscr{D}\xi]$ (the solid black line).

When the news surprise is "small," the reaction of the state-price density, $\mathscr{D}\xi$, becomes a well-defined mathematical object known as a Malliavin derivative.[14] It has the economic meaning of an impulse-response function following a shock in initial values (in our case, a news shock). However, unlike a standard impulse-response function, a Malliavin derivative takes future uncertainty into account: Agent A does not assume that the world becomes deterministic after the shock has occurred and therefore computes an average response.

To obtain an analytical expression for average future returns in our model, notice that the state-price density, $\xi(\delta, \eta)$, in equation (19) depends on two state variables, namely, dividends and the likelihood of Agent B's

[14] See, for example, Detemple and Zapatero (1991), Detemple, Garcia, and Rindisbacher (2003, 2005), Berrada (2006), and Dumas, Kurshev, and Uppal (2009) for applications of Malliavin calculus in financial economics.

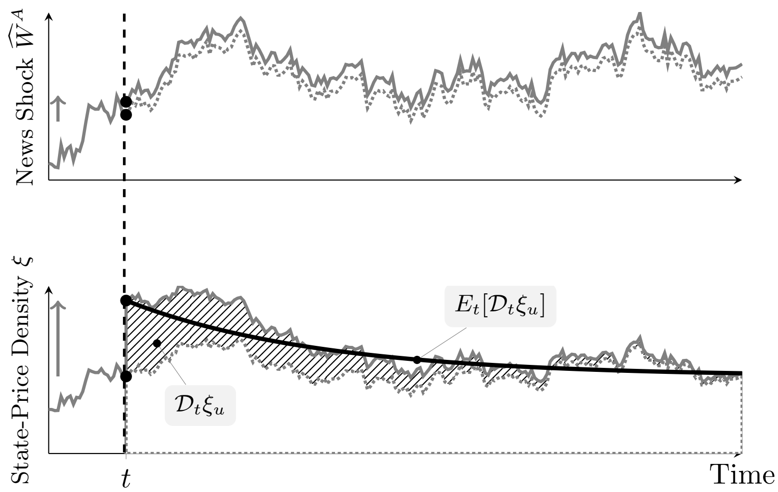

Figure 6. Illustration of the state-price density's response to a news shock. The upper panel shows a trajectory of news before (dashed gray line) and after (solid gray line) a news surprise. The lower panel shows the associated trajectory of the state-price density before (dashed gray line) and after (solid gray line) a news surprise. The shaded area represents the reaction $\mathscr{D}\xi$ to a news shock and the solid black line represents the average reaction $E[\mathscr{D}\xi]$ of the state-price density.

model. As a result, the way the future state-price density reacts to news depends first on how it reacts to a change in these state variables and second on how these state variables themselves react to news. Applying the chain rule to equation (19), the average response of the state-price density satisfies (see Internet Appendix Section VII for derivations):

$$\mathbb{E}_t^{\mathbb{P}^A}\left[R(t,s)\right] = \underbrace{-\alpha \times \left(\sigma_\delta + \frac{\gamma}{\kappa\sigma_\delta}\left(1 - e^{-\kappa(s-t)}\right)\right)}_{\text{fundamental channel } \equiv \frac{\delta_s}{\xi_s}\frac{\partial}{\partial \delta_s}\xi_s \times \frac{\mathscr{D}_t \delta_s}{\delta_s}} + \underbrace{\mathbb{E}_t^{\mathbb{P}^A}\left[(1-\omega_s) \times \left(-\int_t^s \frac{g_u \mathscr{D}_t g_u}{\sigma_\delta^2}\mathrm{d}u\right)\right]}_{\text{disagreement channel } \equiv \mathbb{E}_t^{\mathbb{P}^A}\left[\frac{\eta_s}{\xi_s}\frac{\partial}{\partial \eta_s}\xi_s \times \frac{\mathscr{D}_t \eta_s}{\eta_s}\right]}. \tag{23}$$

We now unravel the chain reaction in equation (23). As news hits, dividends react in a perfectly predictable fashion because the fundamental mean-reverts with constant uncertainty, $\frac{\gamma}{\sigma_\delta}$. The first term in brackets in (23) indicates that dividends respond positively to good news and that their response gradually weakens at speed κ. In contrast, the likelihood of Agent B's model reacts ambiguously to news on average, as news may either corroborate or invalidate Agent B's model *depending on economic conditions*. When the reaction of disagreement $g \times \mathscr{D}g$ is persistently negative, Agent B is overly optimistic and good news thus corroborates her model, as the second term in brackets in (23) shows. When the reaction of disagreement is persistently positive, Agent B is pessimistic and good news invalidates her model.

As dividends and the likelihood of Agent B's model move, the state-price density responds in turn. It decreases after a dividend increase, $\frac{\partial}{\partial \delta}\xi(\delta, \eta) = -\alpha\xi(\delta, \eta)/\delta < 0$, and increases following an increase in the likelihood of Agent B's model, $\frac{\partial}{\partial \eta}\xi(\delta, \eta) = (1 - \omega(\eta))\xi(\delta, \eta)/\eta > 0$. The reason is that an increase in dividends decreases marginal utility. An increase in the likelihood, η, instead increases the risk of Agent A's consumption share growth, which increases her hedging demand and thus her marginal utility, $\phi_A \xi$.

Overall, news shocks move the state-price density through two channels: the fundamental and disagreement. The fundamental channel in (23) triggers a negative state-price density reaction following good news, as it hints future dividend abundance. The disagreement channel leads to an ambiguous reaction depending on whether news corroborates Agent B's model, thereby making Agent A's consumption share growth riskier. Our calibration indicates that the magnitude of the former channel ranges from -0.045 to -0.058, while that of the latter ranges from -0.3 to 0.3. Clearly, disagreement is the main channel driving future returns in our model. As a result, average future returns move together with the average reaction of the likelihood, η, and therefore against the average reaction of disagreement:

$$\operatorname{sign}(\mathbb{E}_t^{\mathbb{P}^A}[R(t,s)]) \approx \operatorname{sign}(\mathbb{E}_t^{\mathbb{P}^A}[\mathscr{D}_t \eta_s]) \equiv -\operatorname{sign}\left(\mathbb{E}_t^{\mathbb{P}^A}\left[\int_t^s g_u \mathscr{D}_t g_u \mathrm{d}u\right]\right). \tag{24}$$

We plot the three-year average future returns in the left panel of Figure 7 along with the average reaction of disagreement, $E_t[g_u \mathscr{D}_t g_u]$, in the right panel. When interpreting average future returns, based on (24) a negative sign implies that returns move opposite to news—they underreact. A positive sign either means that returns adjust to news or overreact, in which case returns subsequently revert—average future returns change direction. Finally, when interpreting the average reaction of disagreement, a negative sign implies that Agent B is overly optimistic while a positive sign means that she is pessimistic.

In bad times (dashed line), news exacerbates disagreement by polarizing agents' opinion. Agent A interprets news positively, whereas Agent B interprets it negatively and disagreement spikes (see right panel). Since news invalidates Agent B's model, returns move opposite to news: good news is followed by negative returns in the short term (see left panel). Hence, Agent B's pessimism slows returns' adjustment to news (decreasing reaction), similar to the underreaction phenomenon in Ottaviani and Sorensen

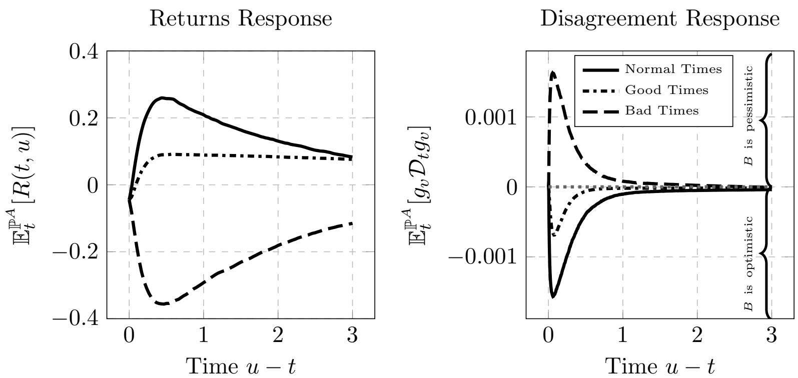

Figure 7. Model-implied impulse response of excess returns to news shocks. The left panel plots the total (sum of the fundamental and disagreement channels) impulse response of stock returns. The right panel plots the response of disagreement to news shocks $\mathbb{E}_t^{\mathbb{P}^A}[g_u \mathscr{D}_t g_u]$. Each line corresponds to a different regime.

(2015).

In normal times (solid line), news shocks precipitate an upward adjustment in Agent B's expectations, sharpening the difference in adjustment speeds across agents. While agents interpret news in similar ways, Agent B overreacts to news and becomes overly optimistic (see the right panel). Because news corroborates Agent B's overoptimism, returns overreact to news in the short term (see the left panel). As a result, Agent B's overreaction accelerates returns' adjustment to news (increasing reaction), similar to the overreaction phenomenon in Daniel, Hirshleifer, and Subrahmanyam (1998).

In good times (dash-dotted line), the reaction of disagreement to news is weak and exhibits little persistence (see the right panel). Because news corroborates both agents' views that economic growth is high, returns adjust immediately (see the left panel). Hence, return continuation arises only when disagreement spikes in the short term. Whether disagreement spikes result from the polarization of opinions in bad times or the difference in adjustment speeds in normal times dictates whether returns under- or over-react. Unlike Ottaviani and Sorensen (2015), under- and overreaction alternate over the business cycle in our model. Finally, because disagreement spikes are countercyclical, news better predicts future returns in bad times.

In the long term (after about six months), excess returns systematically revert, that is, future returns change direction. This later phase of correction arises through the long-term behavior of Agent B's consumption share, $1-\omega$, which multiplies the cumulative disagreement response in equation (23). Since the drift of $1-\omega$,

$$\frac{1}{\mathrm{d}t}\mathbb{E}_t^{\mathbb{P}^A}\left[\mathrm{d}(1-\omega_t)\right] = \frac{g_t^2(2\omega_t - 1 - \alpha)(1-\omega_t)\omega_t}{2\alpha^2\sigma_\delta^2} \leq 0, \quad \forall \alpha \geq 1, \qquad (25)$$

is negative and decreases with the square of disagreement, strong disagreement today leads to a strong downward trend in Agent B's consumption share, which gradually dampens the disagreement response in the right panel of Figure 7. Hence, while a spike in disagreement creates return continuation in the short term, it generates a reversal in the long term.

C. Time-Series Momentum and Momentum Crashes

One of the most pervasive facts in finance is momentum (Jegadeesh and Titman (1993)). Recently, Moskowitz, Ooi, and Pedersen (2012) uncover a similar pattern in aggregate returns, which they coin time-series momentum. In our model, return under- or overreaction to news directly relates to time-series momentum. Disagreement spikes persistently raise the risk premium for holding the stock, thus creating positive serial correlation in excess returns. Disagreement spikes simultaneously reduce Agent B's consumption share over time and thus the risk premium in the future, which ultimately leads to long-term reversal. That disagreement spikes are countercyclical further explains why time-series momentum is stronger in bad times at short horizons and how it crashes following sharp market rebounds.

The economic mechanism for return continuation of Section III.B is based on the concept of returns that matters to investors of the model—the response of the state-price density. However, the same mechanism can be described in terms of time-series momentum based on a common definition of excess returns,

$$r_t^e \equiv \int_{t-\Delta}^{t} \left(\frac{\mathrm{d}S_u + \delta_u \mathrm{d}u}{S_u} - r_u^f \mathrm{d}u \right) = \int_{t-\Delta}^{t} \sigma_u(\theta_u \mathrm{d}u + \mathrm{d}\widehat{W}_u^A), \tag{26}$$

over a period of length Δ. Consider the serial correlation of excess returns at lag h:

$$\rho(h) = \mathrm{cov}^{\mathbb{P}^A}\left(r_{t+h\Delta}^e, r_t^e\right)/\mathrm{var}^{\mathbb{P}^A}\left(r_t^e\right). \tag{27}$$

Following Banerjee, Kaniel, and Kremer (2009), the sign of the coefficient in (27) determines whether returns exhibit momentum ($\rho(h) > 0$) or reversal ($\rho(h) < 0$). Because Brownian innovations in (26) are uncorrelated across nonoverlapping periods, the covariance of instantaneous excess returns, $\sigma\theta$, in (26) dictates the sign of this coefficient.

To show how disagreement spikes affect the coefficient in (27), it is simpler to focus on the covariance of market prices of risk, θ, as opposed to instantaneous excess returns. An approximation of covariances of products (Bohrnstedt and Goldberger (1969)) then yields

$$\begin{aligned}\mathrm{cov}^{\mathbb{P}^A}(\theta_t,\theta_s) \approx\ & \underbrace{\frac{1}{\sigma_\delta^2}\Big(\mathbb{E}^{\mathbb{P}^A}[g_t]\mathbb{E}^{\mathbb{P}^A}[g_s]\mathrm{cov}^{\mathbb{P}^A}(\omega_t,\omega_s)+\mathbb{E}^{\mathbb{P}^A}[1-\omega_t]\mathbb{E}^{\mathbb{P}^A}[1-\omega_s]\mathrm{cov}^{\mathbb{P}^A}(g_t,g_s)\Big)}_{\text{momentum effect } \geq 0} \\ & \underbrace{-\frac{1}{\sigma_\delta^2}\Big(\mathbb{E}^{\mathbb{P}^A}[g_t]\mathbb{E}^{\mathbb{P}^A}[1-\omega_s]\mathrm{cov}^{\mathbb{P}^A}(\omega_t,g_s)+\mathbb{E}^{\mathbb{P}^A}[1-\omega_t]\mathbb{E}^{\mathbb{P}^A}[g_s]\mathrm{cov}^{\mathbb{P}^A}(g_t,\omega_s)\Big)}_{\text{reversal effect } \leq 0}.\end{aligned} \tag{28}$$

Disagreement and agents' consumption share *separately* move the market price of risk in the same direction. In the short term, disagreement keeps moving in the same direction both in bad and in normal times (see Section II), and so does agents' consumption share through (25). Hence, the first term in (28) is positive and induces momentum. In contrast, the *interaction* between disagreement and consumption shares moves current and future prices of risk in opposite directions. In normal and good times, disagreement is negative and moves opposite to Agent A's consumption share; in bad times, it is positive and moves together with Agent A's consumption share. Hence, the second term in (28) is negative and induces reversal.

Whether the momentum effect or the reversal effect in (28) dominates depends on the magnitude of fluctuations in disagreement. Because large and sharp moves in disagreement strengthen the momentum effect in the

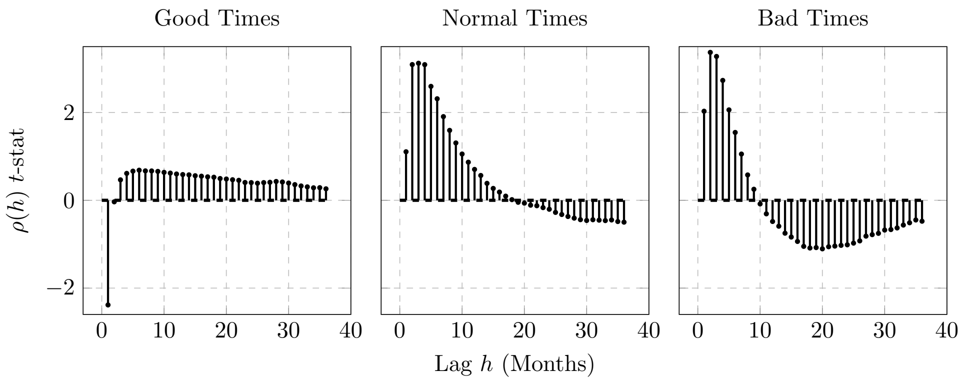

Figure 8. Model-implied conditional time-series momentum. This figure plots the t-statistics of the coefficient $\rho(h)$ for lags h ranging from one month to three years. Each panel corresponds to a different state of the economy. Standard errors are adjusted using Newey and West (1987). The values reported above are obtained from 10,000 simulations of the economy over a 20-year horizon.

short term, disagreement spikes tilt the balance in favor of momentum. To analyze this effect in a way that is consistent with empirical studies, we now consider excess returns at a monthly frequency (i.e., Δ = one month). Figure 8 depicts the t-statistics of the coefficient of serial correlation in equation (27) for lags ranging from one month to three years. The value of the coefficient is reported in Internet Appendix Section XI.B.

The serial correlation of returns has a similar term structure in normal times (center panel) and bad times (right panel): returns exhibit time-series momentum over an horizon of 10 to 18 months, followed by reversal over subsequent horizons. Moreover, the magnitude of momentum varies according to the horizon considered: momentum is large at short horizons (up to four months) and then decays over longer horizons. In Internet Appendix Section XI.C we show that these patterns are also robust unconditionally. Both the magnitude and the timing of time-series momentum are in line with empirical evidence.[15] Indeed, Moskowitz, Ooi, and Pedersen (2012) find significant time series momentum at short horizons (one to six months), weaker momentum at intermediate horizons (seven to 15 months), and reversal at longer horizons.

Since time-series momentum results from a spike in disagreement, the term structures of momentum in normal and bad times differ in an important dimension. The polarization of opinions in bad times induces a sharper spike in disagreement than the difference in adjustment speeds in normal times and thus time-series momentum at short horizons is significantly larger in bad times. That is, more disagreement leads to more time-series momentum. This implication finds strong empirical support (see Section IV). In addition, long-term reversal occurs earlier and is stronger in bad times because a stronger disagreement spike implies a steeper downward trend in Agent B's consumption share (see Section III.B).

Excess returns have different time-series properties in good times (left panel): returns exhibit strong reversal in the very short term and insignificant momentum thereafter. The reason is that news generates little disagreement in good times (see Figure 4). Returns thus immediately revert, as they would if agents had homogeneous beliefs. A consequence of this reversal spike occurring in good times is that a time-series momentum strategy may crash if the market rises sharply. To see this, suppose we implement a momentum strategy in bad or normal times. Figure 8 shows that this strategy remains profitable as long as the economy does not move suddenly

[15] In Internet Appendix Section XI.F we confirm empirically that time-series momentum persists up to one year and is followed by long-term reversal in both NBER expansions and NBER recessions.

to good times. If, instead, the market sharply rebounds, the trend suddenly reverts and the momentum strategy crashes. Sharp trend reversals typically occur at the end of financial crises. For instance, a momentum strategy incurred large losses at the end of the global financial crisis in March, April, and May of 2009 (Moskowitz, Ooi, and Pedersen (2012)).[16] Furthermore, we show in Internet Appendix Section XI.G that the model implies stronger time-series momentum in extreme markets, consistent with Moskowitz, Ooi, and Pedersen (2012).

IV. Testable Predictions and Empirical Evidence

In Section IV.A we construct empirical proxies for the main variables driving excess returns in the model. For instance, we use the dispersion in analysts' forecasts as a proxy for the square of disagreement. We then focus on three novel predictions of the model: *(1)* future excess returns are positively related to contemporaneous dispersion; and time-series momentum at short horizons *(2)* increases with dispersion and *(3)* is strongest in bad times. In Section IV.B we quantify these qualitative predictions using simulations from the model. Repeating the same analysis with observed data, we provide empirical support for these predictions.

A. *Constructing Empirical Proxies for Explanatory Variables*

In the model excess returns depend nonlinearly on three variables—disagreement, the fundamental, and the relative likelihood of agents' models.

[16]Daniel and Moskowitz (2016) and Barroso and Santa-Clara (2015) document a similar phenomenon in the cross-section of stocks.

However, equation (22) shows that this nonlinear dependence is accurately described as a linear combination of two weighted products of these variables, the *weighted fundamental* $\widehat{f}^A(1-\omega(\eta))^2 g^2$ and the *weighted dispersion* $g^2(1-\omega(\eta))^2\omega(\eta)$. We borrow the term "dispersion" from the empirical literature in which heterogeneous beliefs are commonly measured by the dispersion in analysts' forecasts (Diether, Malloy, and Scherbina (2002)), the square of disagreement in our model. Accordingly, we define the monthly weighted fundamental F and the monthly weighted dispersion G as

$$F_t \equiv \int_{t-\Delta}^{t} \widehat{f}_u^A(1-\omega_u)^2 g_u^2 \mathrm{d}u \qquad G_t \equiv \int_{t-\Delta}^{t} g_u^2(1-\omega_u)^2\omega_u \mathrm{d}u. \quad (29)$$

To construct empirical counterparts to F and G in (29), we need empirical proxies for dispersion, the fundamental, and Agent A's consumption share, which we now describe separately. To build an empirical proxy for dispersion, we follow Diether, Malloy, and Scherbina (2002). We first obtain monthly data on analysts' forecasts from I/B/E/S, for the period from February 1976 to November 2013. Denoting by $f^{i,j}$ analyst j's forecast for firm i's fiscal year earnings per share, we define the dispersion D^i in analysts' forecasts about firm i as

$$D_t^i = (|\underset{j}{\mathrm{mean}}(f_t^{i,j})|)^{-1}\underset{j}{\mathrm{std}}(f_t^{i,j}),$$

where $\underset{j}{\mathrm{mean}}(.)$ and $\underset{j}{\mathrm{std}}(.)$ denote the mean and standard deviation of forecasts computed across analysts, respectively. We then compute aggregate dispersion, D, across firms as

$$D_t = \frac{1}{N_t} \sum_{i=1}^{N_t} D_t^i,$$

where N_t is the number of firms in the S&P 500 that have at least two analyst forecasts and a mean forecast different from zero at time t.[17] Regressing the aggregate dispersion D_t on time t shows that dispersion has a strong linear downward trend (see left panel of Table III). Communication and transparency greatly improved over the last 40 years due to technological innovations, for example, the internet and other information technologies. Since our model abstracts from these technological changes, it does not generate a downward trend in the model-implied dispersion. We thus construct a de-trended empirical proxy for dispersion, *Disp*, using the residual from the regression of aggregate dispersion on time.

In Section II we estimate the beliefs $\widehat{f}^A$ and $\widehat{f}^B$ using monthly S&P 500 dividend growth. These beliefs provide a *model-implied* time series of dispersion, $Disp^{MI}$, which is defined as

$$Disp_t^{MI} \equiv g_t^2 = (\widehat{f}_t^A - \widehat{f}_t^B)^2.$$

If the beliefs $\widehat{f}^A$ and $\widehat{f}^B$ offer a reasonable description of observed analyst forecasts, then the model-implied dispersion $Disp^{MI}$ should be correlated with the dispersion among forecasters *Disp*. To investigate this prediction, we regress the observed dispersion, *Disp*, on the model-implied dispersion,

[17]Over the time period 02/1976 to 11/2013, a mean of 490 firms meet these two conditions each month.

Table III

Linear Trend in Dispersion and 12-Month Seasonality in Fundamental

The left and right panels report the outputs obtained by regressing the dispersion D_t on time t and the fundamental GMF on the 12-month seasonality dummy variable Y^{12}, respectively. Standard errors are reported in brackets and statistical significance at the 10%, 5%, and 1% levels is denoted by with * , **, and ***, respectively. Standard errors are adjusted using Newey and West (1987). Data are at the monthly frequency from 02/1976 to 11/2013.

Dep. Var.	D_t
Const.	0.2099***
	(0.0137)
Time t	−0.0032***
	(0.0005)
Adj. R^2	0.1255
Obs.	454

Dep. Var.	GMF_t
Const.	−0.0309***
	(0.0076)
Y_t^{12}	0.1258*
	(0.0660)
Adj. R^2	0.0310
Obs.	454

$Disp^{MI}$. The right panel of Table IV shows that there is a strong positive relation between the two, thus lending support to our assumption that agents use heterogeneous forecasting models. Moreover, conditioning on NBER recessions with the dummy variable Rec, we find that the observed dispersion is higher in recessions than in expansions (left panel of Table IV). The observed dispersion is therefore countercyclical, as our model predicts (see Internet Appendix Section XI.D for an illustration of the time series of $Disp$ and $Disp^{MI}$).

We build a proxy for the fundamental using the simple growth rate, GMF, of the aggregate mean forecast, $MF_t = \frac{1}{N_t}\sum_{i=1}^{N_t} \underset{j}{\text{mean}}(f_t^{i,j})$, which we compute according to

$$GMF_t = (MF_{t-\Delta})^{-1}(MF_t - MF_{t-\Delta}),$$

Table IV

Empirical Dispersion in Recessions and Empirical vs. Model-Implied Dispersion

The left panel reports the outputs obtained by regressing the empirical dispersion, *Disp*, on the dummy variable *Rec*, which equals one during NBER recessions. The right panel reports the outputs obtained by regressing the empirical dispersion, *Disp*, on the model-implied dispersion, $Disp^{MI}$, estimated in Section II. Standard errors are reported in brackets and statistical significance at the 10%, 5%, and 1% levels is denoted by * , **, and ***, respectively. Standard errors are adjusted using Newey and West (1987). Data are at the monthly frequency from 02/1976 to 11/2013.

Dep. Var.	$Disp_t$
Const.	-0.0071
	(0.0044)
Rec_t	0.0525^{***}
	(0.0137)
Adj. R^2	0.0370
Obs.	454

Dep. Var.	$Disp_t$
Const.	-0.0145^{**}
	(0.0060)
$Disp_t^{MI}$	6.6594^{***}
	(1.7589)
Adj. R^2	0.0624
Obs.	454

where $\Delta =$ one month. We do not work with the log growth rate because the aggregate mean forecast is sometimes negative. We also remove three outliers from the time series of growth rates, despite their minor influence on the results. We further eliminate seasonalities from the resulting time series of growth rates by running the regression

$$GMF_t = \alpha + \beta Y_t^{12} + Fund_t,$$

where Y^{12} is a 12-month seasonality dummy variable (see the right panel of Table III). We use the residual of this regression, $Fund$, as an empirical proxy for the fundamental.

Finally, since a model-free proxy for Agent A's consumption share is unavailable, we start by approximating the Brownian shock, $\widehat{W}^A_{t+\Delta} - \widehat{W}^A_t$, as perceived by Agent A using

$$\widehat{W}^A_{t+\Delta} - \widehat{W}^A_t \approx \frac{1}{\sigma_\delta}\left(\log\left(\frac{\delta_{t+\Delta}}{\delta_t}\right) - \left(\widehat{f}^A_t - \frac{1}{2}\sigma^2_\delta\right)\Delta\right), \tag{30}$$

where $\widehat{f}^A_t$ is the belief of Agent A that we estimate in Section II, $\log\left(\frac{\delta_{t+\Delta}}{\delta_t}\right)$ is the S&P 500 dividend growth rate, and $\Delta =$ one month. We then obtain an empirical proxy for the change of measure, which we call *eta*, by substituting the shock in (30) into the discretized dynamics of disagreement and then into the discretized dynamics of the change of measure in equation (8).[18] The empirical proxy for the consumption share of Agent A, which we call *omega*, is obtained by substituting *eta* into equation (12). Figure 9 depicts the historical path of the change of measure, *eta*, and the consumption share of Agent A, *omega*. The consumption share of Agent A varied substantially over the last century, cycling between 0.1 and 0.98 with an average level of 0.61. Although Agent A's share of consumption was larger than that of Agent B on average, both shares of consumption remained mostly within 0.2 and 0.8, with both agents remaining economically influential over the last century.

We now have all the variables we need to compute the empirical coun-

[18]We perform an Euler discretization of the dynamics of disagreement in equation (IA.1).

terparts to the weighted fundamental F and weighted dispersion G in (29):

$$F_{emp} \equiv Fund \times (1 - omega)^2 \times Disp$$
$$G_{emp} \equiv Disp \times (1 - omega)^2 \times omega. \tag{31}$$

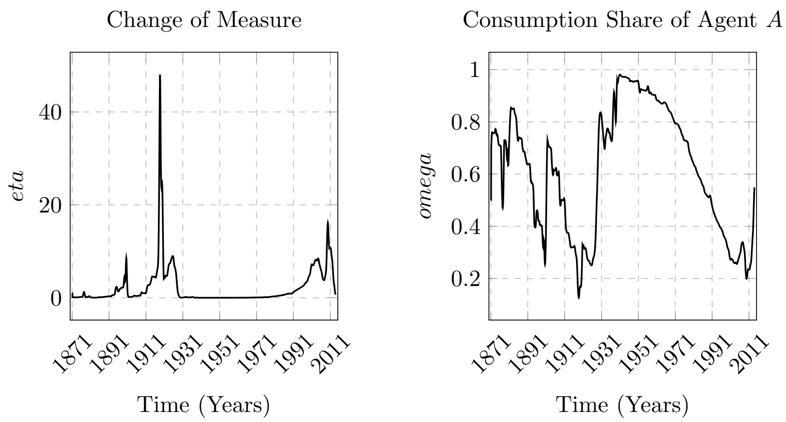

Figure 9. Realized historical path of the change of measure, *eta*, and the corresponding consumption share of Agent *A*, *omega*, from 01/1871 to 11/2013.

B. Testing the Predictions of the Model

We test three new predictions of the model. We show that current dispersion positively predicts future excess returns (Section IV.B.1), and time-series momentum increases with dispersion (Section IV.B.2) and is strongest in bad times (Section IV.B.3).

B.1. Dispersion and Future Excess Returns

To quantify the model-implied relation between excess returns and dispersion at different lags h, we simulate the model and run the regression

$$r^e_{t+h\Delta} = \alpha(h) + \beta_G(h)G_t + \beta_F(h)F_t + \epsilon_{t+h\Delta}, \tag{32}$$

where monthly excess returns r^e are defined in (26). Because weighted dispersion generates most variation in excess returns (see Section III.A), we focus on the coefficient $\beta_G(h)$, which measures the relation between excess returns and dispersion at different lags h. We then run the empirical equivalent to the regression in (32) in which we substitute the weighted fundamental F and the weighted dispersion G by their empirical counterparts F_{emp} and G_{emp} in (31) and take excess returns r^e to be the monthly excess returns on the S&P 500. We plot the t-statistics of the theoretical coefficient $\beta_G(h)$ and its empirical counterpart $\beta_G(h)_{emp}$ in the left and right panels of Figure 10 for lags ranging from one month to one year.

Our model predicts a positive relation between contemporaneous dispersion and future excess returns, the strength of which weakens with the time horizon (see the left panel). To see this, suppose that dispersion is high today. Because dispersion is countercyclical, it will decrease in the future (see Section II.C and Table IV).[19] Since dispersion and excess re-

[19]To confirm that dispersion is countercyclical in our model, we regress the dispersion G_t on the fundamental, $F^m_t \equiv \int_{t-\Delta}^t \widehat{f}^A_u \mathrm{d}u$, using 1,000 simulations over a 100-year horizon. We find that dispersion and the fundamental are indeed negatively related, with a regression coefficient of -0.0017. Using Newey and West (1987) standard errors, the regression coefficient is significant at the 1% level. Note that the regular dispersion, $G_t \equiv \int_{t-\Delta}^t g^2_u \mathrm{d}u$, is also countercyclical.

turns are contemporaneously negatively related (see Section III.A), future excess returns are high, creating a positive relation between contemporaneous dispersion and future excess returns. This observation further implies that future excess returns are countercyclical in our model, consistent with empirical findings. Comparing the two panels of Figure 10 shows that the model can replicate the persistence and the strength of the relation between contemporaneous dispersion and future excess returns that we observe in the data.

Interestingly, the positive relation between *aggregate* dispersion and future *aggregate* returns that we document contrasts with the negative relation that is documented in the cross-section of firms (Diether, Malloy, and Scherbina (2002)). A possible explanation is that short-selling costs, which may create a negative relation between dispersion and future returns (Miller (1977)), matter for the cross-section of firms—costs are high for small companies' stocks—but not in the aggregate market. Consistent with this explanation, Anderson, Ghysels, and Juergens (2005), Boehme et al. (2009), and Carlin, Longstaff, and Matoba (2014) find a positive cross-sectional relation between dispersion and future returns in markets that involve low short-selling costs.

B.2. Time-Series Momentum in High Dispersion Periods

Short-term time-series momentum increases with dispersion in our model. To measure this effect, we define periods of high dispersion with a dummy variable, $Y_{G,t}(p)$, that takes the value of one when the monthly dispersion

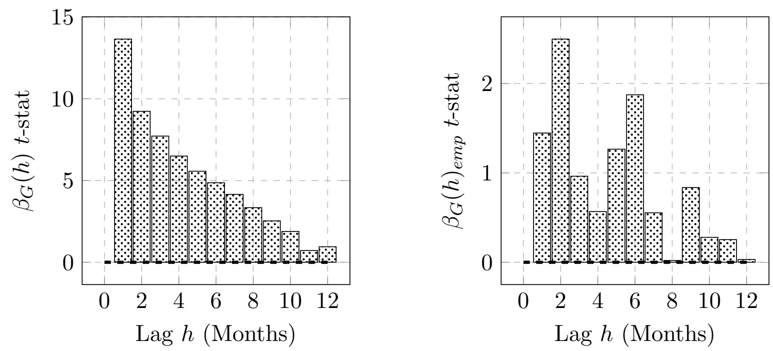

Figure 10. Model-implied and empirical relation between future excess returns and dispersion. The left panel plots the t-statistics of the model-implied response of excess returns to weighted dispersion $\beta_G(h)$ for lags h ranging from one month to one year. The values reported are obtained from 1,000 simulations of the economy over a 100-year horizon. The right panel plots the t-statistics of the empirical response of excess returns to weighted dispersion $\beta_G(h)_{emp}$ for lags h ranging from one month to one year. Data are at the monthly frequency from 02/1976 to 11/2013. All standard errors are adjusted using Newey and West (1987).

G is larger than its p^{th} percentile. We then run the regression

$$r^e_{t+\Delta} = \alpha_G(p) + \beta_{1,G}(p) r^e_t + \beta_{2,G}(p) r^e_t Y_{G,t}(p) + \epsilon_{t+\Delta}, \tag{33}$$

where the model-implied coefficient $\beta_{2,G}(p)$ measures excess time-series momentum in periods of high dispersion. We repeat this exercise by substituting in (33) the weighted dispersion G by its empirical counterpart G_{emp} in (31) and taking excess returns r^e to be the monthly S&P 500 excess returns. The empirical coefficient $\beta_{2,G}(p)_{emp}$ measures excess time-series momentum at a one-month lag in high dispersion periods.[20,21] We plot the t-statistics

[20] We present an alternative approach for estimating the positive relation between time-series momentum and dispersion in Internet Appendix Sections IX and X.

[21] Our results are qualitatively similar if we replace the weighted dispersion $G_t \equiv \int_{t-\Delta}^t g_u^2 (1-\omega_u)^2 \omega_u \mathrm{d}u$ (respectively, $G_{emp} \equiv Disp \times (1-omega)^2 \times omega$) by the regular dispersion $G_t \equiv \int_{t-\Delta}^t g_u^2 \mathrm{d}u$ (respectively, $G_{emp} \equiv Disp$).

of $\beta_{2,G}(p)$ and $\beta_{2,G}(p)_{emp}$ in separate panels in Figure 11 for percentiles ranging from 10 to 50.[22]

Both in the model and in the data, time-series momentum at a one-month lag is larger in periods of high dispersion. However, while the significance of the model-implied relation between dispersion and excess time-series momentum is robust to changes in the dispersion percentile, the empirical relation is strongly significant when the dispersion threshold is the 30^{th} percentile and is relatively weak when the threshold ranges from the 40^{th} to the 50^{th} percentile. The reason for this discrepancy is that spikes in dispersion persist in our model, whereas they revert back quicker in the data (see Internet Appendix Section XI.D). Furthermore, the data also lend support to the prediction that there is short-term time-series reversal in periods of low dispersion (see Internet Appendix Section XI.E).

B.3. Time-Series Momentum over the Business Cycle

The model predicts strongest time-series momentum in bad times at short horizons, a result we quantify as follows. We first identify bad times in our model in a way that can be mapped into our empirical analysis, in which we use NBER recession dates. Since NBER recessions have historically accounted for 30% of the business cycle, we capture recessions with a dummy variable Y_{F^m} that takes the value of one if the fundamental,

[22]We describe the statistics $\beta_{2,G}(p)$ and $\beta_{2,G}(p)_{emp}$ for momentum at a one-month lag only. A positive relation between time-series momentum and dispersion also holds at a two-, three-, and four-month lag.

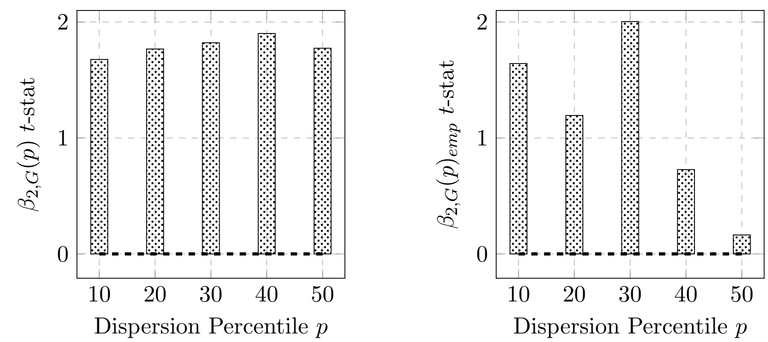

Figure 11. Model-implied and empirical excess one-month time-series momentum in high dispersion periods. The left panel plots the model-implied t-statistics of excess one-month time-series momentum $\beta_{2,G}(p)$ when weighted dispersion is larger than its p^{th} percentile. These values are obtained from 1,000 simulations of the economy over a 100-year horizon. The right panel plots the empirical t-statistics of excess one-month time-series momentum $\beta_{2,G}(p)_{emp}$ when weighted dispersion is larger than its p^{th} percentile. Data are at the monthly frequency from 02/1976 to 11/2013. All standard errors are adjusted using Newey and West (1987).

$F_t^m \equiv \int_{t-\Delta}^{t} \widehat{f}_u^A \mathrm{d}u$, is below its 30^{th} percentile.[23] We then simulate the model and run the regression

$$r_{t+h\Delta}^e = \alpha_F(h) + \beta_{1,F}(h) r_t^e + \beta_{2,F}(h) r_t^e Y_{F^m,t} + \epsilon_{t+h\Delta}, \tag{34}$$

where the model-implied coefficient $\beta_{2,F}(h)$ measures excess time-series momentum in recessions at lag h.

To construct the empirical equivalent to the regression in (34), we specify Y_{F^m} as a dummy variable that takes the value of one during NBER recessions and take excess returns r^e to be monthly S&P 500 excess returns. The empirical coefficient $\beta_2(h)_{emp}$ measures excess time-series momentum in recessions at lag h. To cover sufficiently many business cycle turning

[23]Note that our results are robust to any other threshold lying between the 10^{th} and the 50^{th} percentiles.

points, we extend the sample period and focus on monthly S&P 500 excess returns from January 1871 to November 2013. Looking back to the 19^{th} century allows us to account for a large number of recessions—NBER reports 29 recessions since the beginning of our sample. We plot the t-statistics of $\beta_{2,F}(h)$ and $\beta_2(h)_{emp}$ in separate panels in Figure 12 for lags ranging from one month to one year.

The model accurately replicates excess time-series momentum in recessions and its level of statistical significance at the one-month lag. However, while the model implies persistent excess time-series momentum up to the five-month lag, excess time-series momentum vanishes in the data over lags beyond one month. This discrepancy in persistence between model-implied and observed excess time-series momentum in bad times results from the difference in persistence between model-implied and observed spikes in dispersion. To bring about closer alignment in persistence between the two, disagreement spikes would need to revert faster. For instance, introducing *transient* states in Agent B's Markov chain would increase the reversion speed of her expectations—the sum of all transition intensities—reducing the persistence of disagreement and thus the persistence of excess time-series momentum.

Our finding that short-term time-series momentum is stronger in recessions contrasts with evidence reported in the cross-section of returns: Chordia and Shivakumar (2002) and Cooper, Gutierrez, and Hameed (2004) find that cross-sectional momentum is weaker in downmarkets than in upmarkets. While Moskowitz, Ooi, and Pedersen (2012) show that time-series and cross-sectional momentum are related, their relation seems to vary over the

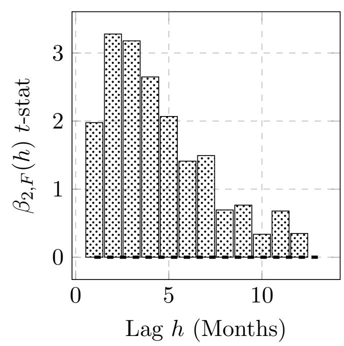

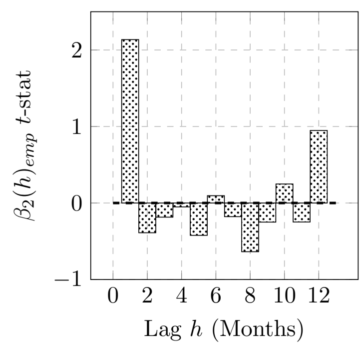

Figure 12. Model-implied and empirical excess time-series momentum in recessions. The left panel plots the t-statistics of the model-implied excess time-series momentum in recessions $\beta_{2,F}(h)$ for lags h ranging from one month to one year. The values reported above are obtained from 1,000 simulations of the economy over a 100-year horizon. The right panel plots the t-statistics of the empirical excess time series momentum in recessions $\beta_2(h)_{emp}$ for lags h ranging from one month to one year. Data are at the monthly frequency from 01/1871 to 11/2013. All standard errors are adjusted using Newey and West (1987).

business cycle.

V. Conclusion

This paper suggests at least two interesting avenues for future research. First, our analysis describes how a single stock—an index—reacts to news shocks, but individual stocks in the index may react differently. In particular, the performance of one stock relative to another may vary over the business cycle. Some may be "losers," while others may be "winners," and this relation may persist or revert depending on economic conditions. Extending our framework to an economy with two trees would allow us to study this cross-sectional relation. Second, in our framework, investors estimate heterogeneous models, both of which measure a different aspect of the business cycle. In a representative-agent economy, we would like to study how the agent dynamically picks one model over the other depending on economic conditions. We believe that such endogenous "paradigm shifts" can explain several empirical facts regarding the dynamics of stock return volatility.

Initial submission: November 19, 2014; Accepted: January 5, 2017

Editors: Bruno Biais, Michael R. Roberts, and Kenneth J. Singleton

REFERENCES

Albuquerque, Rui, and Jianjun Miao, 2014, Advance information and asset prices, *Journal of Economic Theory* 149, 236–275.

Anderson, Evan W., Eric Ghysels, and Jennifer L. Juergens, 2005, Do heterogeneous beliefs matter for asset pricing?, *Review of Financial Studies* 18, 875–924.

Andrei, Daniel, and Julien Cujean, 2017, Information percolation, momentum and reversal, *Journal of Financial Economics* 123, 617–645.

Banerjee, Snehal, Ron Kaniel, and Ilan Kremer, 2009, Price drift as an outcome of differences in higher-order beliefs, *Review of Financial Studies* 22, 3707–3734.

Barberis, Nicholas, Andrei Shleifer, and Robert Vishny, 1998, A model of investor sentiment, *Journal of Financial Economics* 49, 307–343.

Barroso, Pedro, and Pedro Santa-Clara, 2015, Momentum has its moments, *Journal of Financial Economics* 116, 111–120.

Berk, Jonathan B., Richard C. Green, and Vasant Naik, 1999, Optimal investment, growth options, and security returns, *Journal of Finance* 54, 1553–1607.

Berrada, Tony, 2006, Incomplete information, heterogeneity, and asset pricing, *Journal of Financial Econometrics* 4, 136–160.

Biais, Bruno, Peter Bossaerts, and Chester Spatt, 2010, Equilibrium asset pricing and portfolio choice under asymmetric information, *Review of Financial Studies* 23, 1503–1543.

Boehme, Rodney D., Bartley R. Danielsen, Praveen Kumar, and Sorin M. Sorescu, 2009, Idiosyncratic risk and the cross-section of stock returns: Merton (1987) meets miller (1977), *Journal of Financial Markets* 12, 438–468.

Bohrnstedt, George W., and Arthur S. Goldberger, 1969, On the exact covariance of products of random variables, *Journal of the American Statistical Association* 64, 1439–1442.

Bollerslev, Tim, and Robert J. Hodrick, 1992, *Financial Market Efficiency Tests* (National Bureau of Economic Research).

Brennan, Michael J., and Yihong Xia, 2001, Stock price volatility and equity premium, *Journal of Monetary Economics* 47, 249–283.

Buraschi, Andrea, Fabio Trojani, and Andrea Vedolin, 2014, When uncertainty blows in the orchard: Comovement and equilibrium volatility risk premia, *Journal of Finance* 69, 101–137.

Buraschi, Andrea, and Paul Whelan, 2013, Term structure models and differences in beliefs, Working paper, Imperial College.

Carlin, Bruce I., Francis A. Longstaff, and Kyle Matoba, 2014, Disagreement and asset prices, *Journal of Financial Economics* 114, 226–238.

Cen, Ling, K. C. John Wei, and Liyan Yang, 2017, Disagreement, underreaction, and stock returns, *Management Science* 63, 1214–1231.

Cespa, Giovanni, and Xavier Vives, 2012, Dynamic trading and asset prices: Keynes vs. Hayek, *Review of Economic Studies* 79, 539–580.

Chalkley, Martin, and In Ho Lee, 1998, Learning and asymmetric business cycles, *Review of Economic Dynamics* 1, 623–645.

Chordia, Tarun, and Lakshmanan Shivakumar, 2002, Momentum, business cycle, and time-varying expected returns, *Journal of Finance* 57, 985–1019.

Cooper, Michael J., Roberto C. Gutierrez, and Allaudeen Hameed, 2004, Market states and momentum, *Journal of Finance* 59, 1345–1365.

Cox, John C., and Chi-fu Huang, 1989, Optimal consumption and portfolio policies when asset prices follow a diffusion process, *Journal of Economic Theory* 49, 33–83.

Dangl, Thomas, and Michael Halling, 2012, Predictive regressions with time-varying coefficients, *Journal of Financial Economics* 106, 157–181.

Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 1998, Investor psychology and security market under- and overreactions, *Journal of Finance* 53, 1839–1885.

Daniel, Kent, and Tobias J. Moskowitz, 2016, Momentum crashes, *Journal of Financial Economics* 122, 221–247.

David, Alexander, 1997, Fluctuating confidence in stock markets: Implications for returns and volatility, *Journal of Financial and Quantitative Analysis* 32, 427–462.

David, Alexander, 2008, Heterogeneous beliefs, speculation, and the equity premium, *Journal of Finance* 63, 41–83.

Detemple, Jerome B., 1986, Asset pricing in a production economy with incomplete information, *Journal of Finance* 41, 383–391.

Detemple, Jerome B., 1991, Further results on asset pricing with incomplete information, *Journal of Economic Dynamics and Control* 15, 425–453.

Detemple, Jerome B., René Garcia, and Marcel Rindisbacher, 2003, A Monte Carlo method for optimal portfolios, *Journal of Finance* 58, 401–446.

Detemple, Jerome B., René Garcia, and Marcel Rindisbacher, 2005, Representation formulas for Malliavin derivatives of diffusion processes, *Finance and Stochastics* 9, 349–367.

Detemple, Jerome B., and Shashidhar Murthy, 1994, Intertemporal asset pricing with heterogeneous beliefs, *Journal of Economic Theory* 62, 294–320.

Detemple, Jerome B., and Fernando Zapatero, 1991, Asset prices in an exchange economy with habit formation, *Econometrica* 59, 1633–1657.

Diether, Karl B., Christopher J. Malloy, and Anna Scherbina, 2002, Differences of opinion and the cross section of stock returns, *Journal of Finance* 57, 2113–2141.

Dumas, Bernard, Alexander Kurshev, and Raman Uppal, 2009, Equilibrium portfolio strategies in the presence of sentiment risk and excess volatility, *Journal of Finance* 64, 579–629.

Ehling, Paul, Michael Gallmeyer, Christian Heyerdahl-Larsen, and Philipp Illeditsch, 2017, Disagreement about inflation and the yield curve, *Forthcoming, Journal of Financial Economics.*

Garcia, Diego, 2013, Sentiment during recessions, *Journal of Finance* 68, 1267–1300.

Grinold, Richard, and Ronald Kahn, 1999, *Active Portfolio Management: A Quantitative Approach for Producing Superior Returns and Controlling Risk* (McGraw-Hill Education).

Hamilton, James D., 1989, A new approach to the economic analysis of nonstationary time series and the business cycle, *Econometrica* 57, 357–384.

Hamilton, James D., 1994, *Time Series Analysis* (Princeton University Press).

Henkel, Sam James, J. Spencer Martin, and Federico Nardari, 2011, Time-varying short-horizon predictability, *Journal of Financial Economics* 99, 560–580.

Hodrick, Robert J., and Edward C. Prescott, 1997, Postwar U.S. business cycles: An empirical investigation, *Journal of Money, Credit and Banking* 29, 1–16.

Holden, Craig W., and Avanidhar Subrahmanyam, 2002, News events, information acquisition, and serial correlation, *Journal of Business* 75, 1–32.

Hong, Harrison, and Jeremy C. Stein, 1999, A unified theory of underreaction, momentum trading, and overreaction in asset markets, *Journal of Finance* 54, 2143–2184.

Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48, 65–91.

Johnson, Timothy C., 2002, Rational momentum effects, *Journal of Finance* 57, 585–608.

Kandel, Eugene, and Neil D. Pearson, 1995, Differential interpretation of public signals and trade in speculative markets, *Journal of Political Economy* 103, 831–872.

Kogan, Leonid, Stephen A. Ross, Jiang Wang, and Mark M. Westerfield, 2006, The price impact and survival of irrational traders, *Journal of Finance* 61, 195–229.

Lipster, Robert S., and Albert N. Shiryaev, 2001, *Statistics of Random Processes II* (Springer-Verlag, New York).

Loh, Roger K., and René M. Stulz, 2014, Is sell-side research more valuable in bad times? Working paper, Ohio State University.

Makarov, Igor, and Oleg Rytchkov, 2012, Forecasting the forecasts of others: Implications for asset pricing, *Journal of Economic Theory* 147, 941–966.

Milas, Costas, Philip Rothman, and Dick van Dijk, 2006, *Nonlinear Time Series Analysis of Business Cycles* (Emerald Group Publishing).

Miller, Edward M., 1977, Risk, uncertainty, and divergence of opinion, *Journal of Finance* 32, 1151–1168.

Moskowitz, Tobias J., Yao Hua Ooi, and Lasse Heje Pedersen, 2012, Time series momentum, *Journal of Financial Economics* 104, 228–250.

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Ottaviani, Marco, and Peter Norman Sorensen, 2015, Price reaction to information with heterogeneous beliefs and wealth effects: Underreaction, momentum, and reversal, *American Economic Review* 105, 1–34.

Patton, Andrew, and Allan Timmermann, 2010, Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion, *Journal of Monetary Economics* 7, 803–820.

Piatti, Ilaria, and Fabio Trojani, 2015, Predictable risks and predictive regression in present-value models, Working paper, Swiss Finance Institute.

Rapach, David E., Jack K. Strauss, and Guofu Zhou, 2010, Out-of-sample equity premium prediction: Combination forecasts and links to the real economy, *Review of Financial Studies* 23, 821–862.

Sagi, Jacob S., and Mark S. Seasholes, 2007, Firm-specific attributes and the cross-section of momentum, *Journal of Financial Economics* 84, 389–434.

Scheinkman, José A., and Wei Xiong, 2003, Overconfidence and speculative bubbles, *Journal of Political Economy* 111, 1183–1219.

Sims, Christopher A., 1980, Macroeconomics and reality, *Econometrica* 48, 1–48.

Stewart, Scott, Christopher Piros, and Jeffrey Heisler, 2010, *Running Money: Professional Portfolio Management* (McGraw-Hill/Irwin).

Tetlock, Paul C., 2007, Giving content to investor sentiment: The role of media in the stock market, *Journal of Finance* 62, 1139–1168.

Van Nieuwerburgh, Stijn, and Laura Veldkamp, 2006, Learning asymmetries in real business cycles, *Journal of Monetary Economics* 53, 753–772.

Vayanos, Dimitri, and Paul Woolley, 2013, An institutional theory of momentum and reversal, *Review of Financial Studies* 26, 1087–1145.

Veldkamp, Laura L., 2005, Slow boom, sudden crash, *Journal of Economic Theory* 124, 230–257.

Veronesi, Pietro, 1999, Stock market overreaction to bad news in good times: A rational expectations equilibrium model, *Review of Financial Studies* 12, 975–1007.

Veronesi, Pietro, 2000, How does information quality affect stock returns?, *Journal of Finance* 55, 807–837.

Vuong, Quang H., 1989, Likelihood ratio tests for model selection and non-nested hypotheses, *Econometrica* 57, 307–333.

Xiong, Wei, 2014, *Bubbles, Crises, and Heterogeneous Beliefs* (Handbook for Systemic Risks).

Xiong, Wei, and Hongjun Yan, 2010, Heterogeneous expectations and bond markets, *Review of Financial Studies* 23, 1433–1466.