Factor and stock-specific disagreement and trading flows[*]

Fotis Grigoris Christian Heyerdahl-Larsen Preetesh Kantak

April 2022

Abstract

We study how disagreement on both factor and stock-specific risk exposures across many agents and securities impacts asset prices. Our theoretical analyses predict that disagreement about factor dynamics drives larger flows into portfolios that are more exposed to the factors. Consequently, these concentrated bets on the factor lead to higher volatility and reduced diversification benefits. We then test these predictions using a novel empirical setting – exchange-traded funds (ETFs). We find that when factor disagreement rises, funds flow into the ETFs that mimic the factor. However, these increased flows induce high forward-looking volatility of, and correlation risk within, the ETF.

1 Introduction

A growing literature on crowded trades focuses on concentration in individual stock positions and the realization of tail risks during crisis events (inter alia, Brown et al. (2021), Ben-David et al. (2012), and Ellul et al. (2011)). This paper, in contrast, considers how disagreement in subjective expectations about the stock-specific and factor components of expected returns (i) induces trading activity that leads to crowded trades, and (ii) is priced by financial markets. On the theoretical front, we consider *why* these concentrated positions are created (flow generated from time-varying factor disagreement) and *how* the risks are embedded in asset prices (higher volatility of and correlation between stocks that compose the factor). On the empirical front, we quantify the

[*]We thank Fahiz Baba-Yara, Gurdip Bakshi, Matteo Binfarè, Gregory Brown, Riccardo Colacito, Chris Lundblad, Noah Stoffman, Chuck Trzcinka, Motohiro Yogo, and seminar participants at Indiana University and St. Louis Federal Reserve for helpful comments and suggestions. Grigoris (fgrigori@iu.edu), Heyerdahl-Larsen (chheyer@iu.edu), and Kantak (pkantak@iu.edu@iu.edu) are at the Kelley School of Business at Indiana University.

effects of disagreement-induced trading implied by our theory within a relatively unique setting: exchange-traded funds (ETFs). Given the rapid growth of trading in ETF-linked options over the last decade, ETFs provide a rich empirical setting to examine how disagreement impacts ex-ante and long-run perceptions, rather than realizations, of risk.

The canonical asset pricing framework assumes that variations in the returns of individual securities are a function of systematic risk factor exposures and stock-specific shocks. This sets up a dichotomy: with many investors and many stocks there are a near infinite number of dimensions along which investors can disagree on the stock-specific portion of returns. However, this form of disagreement (henceforth, *stock-specific disagreement*) should wash out in the aggregate. In contrast, investors can disagree on far fewer dimensions regarding the portion of returns attributed to systematic risk factors (henceforth, *factor disagreement*). This generates the distinct possibility of group-think and concentrated exposure (both long and short) across investors. One would think that each form of disagreement would have a different impact on asset prices and different consequences for financial risk.

To demonstrate why, consider a situation in which investors are considering trading on their beliefs about the prospects of the U.S. technology sector. On the one hand, if investors' beliefs are *dispersed*, meaning that investors believe that stock-specific shocks to firms are likely to drive next period's returns, then investors will choose to take disparate positions in individual stocks, such as Microsoft Corporation (NASDAQ: MSFT) and Apple Corporation (NASDAQ: AAPL). On the other hand, if investors believe that the returns of technology firms will be driven by these firms' exposures to a common systematic "technology" factor, then investors will likely choose a portfolio of technology stocks. For instance, investors may trade the Invesco QQQ Trust Series 1 ETF (NASDAQ: QQQ), an ETF that tracks the returns of the NASDAQ-100 index. Consequently, a large inflow into (outflow from) QQQ suggests that investors are predominately trading on the systematic (firm-specfic) component of technology stock returns.

Our analysis begins by constructing an economic model that codifies this intuition. Specifically, we consider a pure exchange economy with multiple Lucas trees that are exposed to both factor risk and stock-specific shocks. The model's primary innovation is that investors can disagree about both dimensions of returns. In our model, there are periods when there is strong disagreement about the factor, and periods when there is strong disagreement about stock-specific returns (i.e., dispersed disagreement). Although there is still a lot of disagreement in times of dispersed disagreement, different agents take (uncorrelated) bets on different stocks, and hence the effect on the aggregate portfolio is muted. In contrast, factor disagreement drives all investors to take correlated bets on the systematic component of returns, inducing a large impact on the portfolios that mimic the factors, which are represented as ETFs in the data.

The key model predictions are fourfold: first, factor disagreement increases the exposure of investors in the economy to the factor. In the context of our empirical analysis, this translates into greater flow into instruments that are primarily exposed to factor risk (i.e., ETFs). Second, factor disagreement increases the return volatility of ETFs, as these instruments closely align with the systematic risk factor. Most critically, this increase in volatility is caused by the higher correlations between stocks that compose the factor, translating into consequences for factor exposure risk. Finally, there is a U-shaped relationship between factor disagreement and the volatility of the individual securities composing the factor. When factor disagreement is a small (large) fraction of total disagreement, stock-specific (factor) disagreement dominates investor trading incentives, leading to higher individual security volatility.

We use the return and flow dynamics of ETFs and their underlying securities to test these hypotheses. ETFs provide a novel and interesting crucible for our tests for three primary reasons. First, it is difficult to empirically analyze how dynamics in disagreement impact stock prices directly because of a missing variable: flow. In any theoretical framework, changes in disagreement induces changes in trading between agents. It is these trading dynamics, not those of disagreement itself, that drive changes in the risks and returns of assets. In contrast, most papers in this area of research study an *inferred* association, testing how measures of disagreement ultimately predict returns or changes in covariances (see, e.g., Buraschi et al. (2014) and Daniel et al. (2021)) without considering the first-order effects on trade flows.

ETFs allow us to circumvent the aforementioned shortcomings and examine how differences in disagreement both across ETFs and over time directly impact trading activity, and consequently affect asset prices. As ETFs track specific indices, and investment "themes" and "styles," authorized participants are incentivized to maintain the connection between the underlying asset value of the ETF and that of its component securities (see, e.g., Ben-David et al. (2018) for details). If there is even an infinitesimally higher cost of trading a basket of securities than directly trading an index itself, an increase (decrease) in aggregate demand for diversified factor exposure should lead to the creation (destruction) of ETFs. In our first piece of analysis, we exploit this connection to show how changes in factor or stock-specific disagreement drive changes in the creation or destruction of ETFs. In particular, we find that a one-standard-deviation higher amount of factor disagreement measure (relative to the stock-specific disagreement measure) leads to a 0.08 standard-deviation higher flow into the ETF. This positive association between the relative amount of factor vis-á-vis stock-specific disagreement and trade flows is robust to controlling for both ETF and time fixed effects, as well as a variety of confounding variables that could drive ETF flows (e.g., lagged ETF returns). Altogether, the positive relation between the importance of factor disagreement and trading flows is in line with the theoretical predictions of our model.

Second, we are the first to our knowledge to exploit the richness of the ETF options market, which today composes more than 40% of all option volume, for our analysis.[1]. Most papers that link ETF trading activity to risk use the physical return space for their analyses. By using ETF option prices, we can link disagreement and flow directly to changes in investor perceptions, rather than noisy realizations, of risk. Examining risk-neutral moments allows us to more accurately measure and identify forward-looking and, more importantly, long-dated shifts in the volatility of, and correlation between, securities exposed to a given factor. Finally, ETFs provide us with a tremendous cross section of various investment factors and styles in which to analyze our model's novel predictions. We document a significant amount of heterogeneity in disagreement, volatility, and correlations across ETFs.

In keeping with our theoretical analysis, we also find that risk-neutral (or forward-looking) ETF volatility is strongly related to the relative amount of disagreement about common (i.e., factor or systematic) component of an ETF's returns. A one-standard-deviation increase in our relative factor disagreement measure leads to volatility rising by about 0.06 standard deviations, even when accounting for ETF and time fixed effects and a battery of ETF-level controls.

Finally, as the volatility of individual cash flows in our model is constant, the relation between flow into an ETF and the ETF's volatility is driven by changes in the correlations between securities. Consequently, flows into an ETF impact the correlation risk between the stocks composing the

[1] See, for example, the Wall Street Journal article "Just as Hot as ETFs: Options on ETFs" from December 9, 2019.

ETF, and alter the diversification benefits of investing in the index. To show this, we compute the average risk-neutral correlation between all pairs of securities within an ETF and find that increased disagreement about systematic risk exposure of an ETF is indeed strongly and positively related to the average intra-ETF correlation. When the relatively amount of systematic disagreement rises by one standard deviation, then the average risk-neutral correlation between all pairs of securities with an ETF rises by about 0.03 standard deviations. This indicates that increased systematic disagreement about the drivers of an ETF's returns are associated with a *decline* in the diversification benefits of holding the ETF, as a result of the anticipated increase in flows into the ETF. Moreover, we find that this loss of diversification benefits is concentrated among one- to six-month ahead estimates of correlation. This contrasts with other papers that focus on losses in diversification benefits over one month (see, e.g., Da and Shive (2018)).

Contributions and related literature. Many have used the arbitrage relationship between ETFs and their underlying securities for economic analysis. Ben-David et al. (2018) analyzes how the level of ETF ownership predicts single security volatility and mispricing, and Da and Shive (2018) analyze how ETF trading activity induces changes in physical correlation. Most recently, Huang et al. (2021) demonstrate how informed trading improves the market efficiency of the component securities. Unlike these papers, we use ETFs and their options to analyze the dynamic effects of *disagreement* on ETF and individual security flows, volatility, and correlations. Option-based measures are uniquely suited to speak to the forward-looking and long-dated implications of this activity. The previous literature finds primarily short-run and mean-reverting effects, which suggests that mispricing drive the results. In contrast, our results are primarily driven by risk premia. Our analysis and findings are therefore also closely aligned with the work on indexation (see, e.g., Barberis et al. (2005), Baltussen et al. (2019), Brogaard et al. (2019), and Bond and Garcia (2021)).

Our paper is also related to the literature on general equilibrium models with heterogeneous beliefs such as Harrison and Kreps (1978), Detemple and Murthy (1994), Zapatero (1998), and Basak (2000), among others.[2] With some exceptions, most of the literature on disagreement has

[2]Models with disagreement include, among many others, Scheinkman and Xiong (2003), Basak (2005), Berrada (2006), Buraschi and Jiltsov (2006), Jouini and Napp (2007), David (2008), Dumas et al. (2009), Xiong and Yan (2010), Cvitanic and Malamud (2011), Cvitanic et al. (2012), Bhamra and Uppal (2014), Buraschi et al. (2014), Cujean and Hasler (2017), Ehling et al. (2018), and Atmaz and Basak (2018).

focused on economies with a single stock or Lucas tree. In contrast, we consider an economy featuring multiple Lucas trees with a factor structure to study the implications of disagreement on factor versus stock-specific risk.

Finally, our paper also relates to the literature on the empirical relations between disagreement and asset prices (see, e.g., Berkman et al. (2009); Chen et al. (2002); Diether et al. (2002); Park (2005)) and disagreement and trading activity (see, e.g., Ajinkya et al. (1991); Bessembinder et al. (1996); Goetzmann and Massa (2005)). In the spirit of this literature, we construct measures of disagreement using the dispersion of analysts' forecasts and earnings estimates. However, unlike many of the aforementioned studies, we use these analysts estimates to construct measures of disagreement that are specific to each ETF in our sample. Our ETF-specific measures of disagreement capture both the extent to which brokers disagree about the common factor underlying an ETF and the extent to which analysts disagree about the idiosyncratic component of an ETF's earnings. These two proxies map to our notions of factor and stock-specific disagreement, respectively.

2 Model

We consider a continuous time model with incomplete information, multiple agents with heterogeneous beliefs, and multiple Lucas trees.

2.1 Output

We assume that there are $N+1$ Lucas trees in the economy. The dynamics of the first tree is

$$dE_t = \mu_E E_t dt + \sigma_E E_t dw_{E,t}, \tag{1}$$

where $w_{E,t}$ is a standard Brownian motion independent of all other shocks in the model. The first Lucas tree, E_t, is assumed to represent any endowment that is not from dividends (e.g., labor). The remaining trees represent dividends from stocks. We assume that the dividend at time t for stock $n = 1, \dots, N$ is

$$D_{n,t} = D_{n,0} E_t e^{z_t + \epsilon_{n,t}}. \tag{2}$$

In Equation (2), E_t, represents the common component of the stock market with the rest of the economy, z_t is an aggregate factor that all stocks load on, and ϵ_t represents an idiosyncratic component. Total dividends are $D_t = \sum_{n=1}^{N} D_{n,t}$ and aggregate consumption is $C_t = E_t + D_t$.

2.2 Information

We assume that agents have full information about the process E_t. However, agents can only observe z_t and $\epsilon_{n,t}$. Specifically, we assume that the true dynamics of the factor, z_t, is

$$dz_t = \mu_z dt + \sigma_z dw_{z,t} \tag{3}$$

and the dynamics of the idiosyncratic components $\epsilon_{n,t}$ are

$$d\epsilon_{n,t} = \mu_n dt + \sigma_n dw_{n,t}, \tag{4}$$

for $n = 1, \ldots, N$ where $w_{z,t}$ and $w_{n,t}$ are standard Brownian motion that are independent of each other. Agents disagree about the dynamics of the factor, z_t, and the idiosyncratic components. Specifically, we assume that agent $j = 1, \ldots, J$ believes that the dynamics is

$$dz_t = \mu^j_{z,t} dt + \sigma_z dw^j_{z,t}, \tag{5}$$

where $\mu^j_{z,t} = \mu_z + \sigma_z \Delta^j_z s_t$. Since all agents observe z_t, this implies that the observed shock $w^j_{z,t}$ is related to the true shock $w_{z,t}$ by

$$dw^j_{z,t} = dw_{z,t} - \Delta^j_z s_t dt. \tag{6}$$

Similarly, we assume that agent j believes that the dynamics of the idiosyncratic component of stock n is

$$d\epsilon_{n,t} = \mu^j_{n,t} dt + \sigma_n dw^j_{n,t}, \tag{7}$$

where $\mu^j_{n,t} = \mu_n + \sigma_n \Delta^j_n \left(1 - s_t\right)$, implying

$$dw^j_{n,t} = dw_{n,t} - \Delta^j_n (1 - s_t) dt. \tag{8}$$

In the above, s_t is a process that governs the composition of the disagreement in the economy. Specifically, we assume that

$$s_t = \frac{1}{1+e^{-\delta_t}}, \quad \text{where} \;\; d\delta_t = \kappa_\delta \left(\bar{\delta} - \delta_t\right) dt + \sigma_\delta dw_{\delta,t}, \tag{9}$$

and $w_{\delta,t}$ is a standard Brownian motion independent of all other shocks. Since s_t is bounded between zero and and one we can interpret s_t as the amount of systematic versus idiosyncratic disagreement in the economy. As s_t approaches one, there is less disagreement about the dynamics of the idiosyncratic processes $\epsilon_{n,t}$ and more disagreement about the factor z_t.

We make an additional symmetry assumption about the beliefs. Specifically, we assume that there are $J = 2N$ agents in the economy. For agent $j = 1, \ldots, N$, we assume that the belief about the systematic component is $\Delta_z^j = \Delta > 0$. This implies that these agents are *optimistic* about the factor growth rate. For the other half of the population we assume that $\Delta_z^j = -\Delta$, implying these agents are *pessimists*. In addition, we assume that the belief about the idiosyncratic component of stock n are as follows

$$\Delta_n^{2n-1} = \Delta \tag{10}$$

$$\Delta_n^{2n} = -\Delta \tag{11}$$

and zero for all other agents. Hence, each agent is associated with a specific stock and is either optimistic or pessimistic about that stock, while remaining neutral for all other stocks. Given that half of the agents are optimistic and half are pessimistic about the factor, we label times with a high s_t as times of high *factor disagreement*. Similarly, in times when s_t is low, there is mostly disagreement about the idiosyncratic component and we label this as times of *stock-specific disagreement*. We stack all the Brownian shocks into a vector $w_t = (w_{E,t}, w_{\delta,t}, w_{z,t}, w_{1,t}, \ldots, w_{N,t}) \in \mathbb{R}^{N+3}$. Similarly, define agent j's disagreement vector at time t as $\Delta_t^j = \Big(0, 0, \Delta_z^j s_t, \Delta_1^j (1-s_t), \ldots, \Delta_N^j (1-s_t)\Big)$. The disagreement vector of agent j captures how distorted the belief is about each shock in the Brownian vector w_t at time t.

2.3 Security markets

Agents can trade a locally risk-free asset in zero net supply, $N+1$ stocks, and two additional zero net supply derivatives. Since the market has $N+3$ shocks, the market is potentially complete. The locally risk-free asset follows

$$dB_t = r_t B_t dt, \tag{12}$$

where r_t is real short rate determined in equilibrium. In addition, since stock $n = 1, \ldots, N$ is a claim to $D_{n,t}$, the return dynamics of stock n is

$$dR_{n,t} = \frac{dS_{n,t} + D_{n,t}dt}{S_{n,t}} = \mu_{R_n,t}dt + \sigma'_{R_n,t}dw_t, \tag{13}$$

where $\sigma_{R_n,t} \in \mathbb{R}^{N+3}$. We solve for $\sigma_{R_n,t}$ and $\mu_{R_n,t}$ in equilibrium. We also assume that agents can trade a claim on the first Lucas tree, E_t, with return dynamics given by

$$dR_{E,t} = \frac{dS_{E,t} + E_t dt}{S_{E,t}} = \mu_{R_E,t}dt + \sigma'_{R_E,t}dw_t. \tag{14}$$

Besides the risk-free asset and the $N+1$ stocks, we assume that there are two zero net supply claims (derivatives) that agents can trade. The first derivative is linked to the shock to δ_t

$$dR_{w_\delta,t} = \mu_{w_\delta,t}dt + dw_{\delta,t}, \tag{15}$$

and the second derivative is linked to the shock to E_t,

$$dR_{w_E,t} = \mu_{w_E,t}dt + dw_{E,t}, \tag{16}$$

where $\mu_{w_\delta,t}$ and $\mu_{w_E,t}$ are determined in equilibrium. It is convenient to summarize the price system in terms of the stochastic discount factor. In our economy, agents have different beliefs, and therefore perceive different market prices of risk. Consequently, each agent perceives the stochastic discount factor differently. The dynamics of the stochastic discount factor as perceived by agent j is

$$dM_t^j = -r_t M_t^j dt - \theta_{E,t} M_t^j dw_{E,t} - \theta_{z,t}^j M_t^j dw_{z,t}^j - \sum_{n=1}^{N} \theta_{n,t}^j M_t^j dw_{n,t}^j. \tag{17}$$

Under the true measure the stochastic discount factor has the dynamics

$$dM_t = -r_t M_t^j dt - \theta_{E,t} M_t dw_{E,t} - \theta_{z,t}^j M_t dw_{z,t} - \sum_{n=1}^{N} \theta_{n,t} M_t dw_{n,t}, \tag{18}$$

where $\theta_{z,t}^j = \theta_{z,t} + \Delta_z^j s_t$ and $\theta_{n,t}^j = \theta_{n,t} + \Delta_n^j (1 - s_t)$. In equilibrium, we have $\mu_{i,t} = r_t + \theta_t' \sigma_{i,t}$ for $i = E, w_E, w_\delta, R_1, \ldots, R_N$. Hence, the expected return perceived by agent j is related to the expected return under the true measure by

$$\mu_{i,t}^j = \mu_{i,t} + \Delta_{z,t}^j s_t \sigma_{i,z,t} + \sum_{n=1}^{N} \Delta_{n,t}^j (1 - s_t) \sigma_{i,n,t}. \tag{19}$$

As noted above, if Δ_z^j (or Δ_n^j) is positive it implies that agent j is optimistic about z (or ϵ_n). In that case, we see from Equation (19) that the agent will also perceive a higher expected return provided that the loading of the asset $\sigma_{i,z,t}$ (or $\sigma_{i,n,t}$) is positive. Note that we can define the disagreement process of agent j by η_t^j that links the perceived stochastic discount factor of agent j, M_t^j, to the stochastic discount factor under the true measure, M_t, by $M_t^j = M_t / \eta_t^j$. The disagreement process is formally a Radon-Nikodym derivative with dynamics

$$d\eta_t^j = \Delta_z^j s_t \eta_t^j dw_{z,t} + \sum_{n=1}^{N} \Delta_n^j (1 - s_t) \eta_t^j dw_{n,t}. \tag{20}$$

2.4 Preferences

Agents maximize lifetime utility given by

$$E^j \left[\int_0^\infty e^{-\rho t} \frac{c_{j,t}^{1-\gamma}}{1-\gamma} dt \right], \tag{21}$$

subject to the dynamic budget constraint

$$\begin{aligned} dW_t^j &= \left(r_t W_t^j + \pi_{E,t}^j \left(\mu_{E,t}^j - r_t\right) + \pi_{w_E,t}^j \left(\mu_{w_E,t} - r_t\right)\right) dt \\ &+ \left(\pi_{w_\delta,t}^j \left(\mu_{w_\delta,t} - r_t\right) + \sum_{n=1}^{N} \pi_{n,t}^j \left(\mu_{R_n,t}^j - r_t\right) - c_{j,t}\right) dt \\ &+ \pi_{w_E,t}^j dw_{E,t} + \pi_{w_\delta,t}^j dw_{\delta,t} \\ &+ \pi_{E,t}^j \sigma_{E,t}' dw_t + \sum_{n=1}^{N} \pi_{n,t}^j \sigma_{R_n,t}' dw_t, \end{aligned} \tag{22}$$

with $W_0^j = w^j$ and where $\pi_{i,t}^j$ for $i = E, w_E, w_\delta, 1, \ldots, N$ is the dollar amount invested in asset i by agent j. Note the the expectation in Equation (21) and the dynamics of the wealth in Equation (22) are under the belief of agent j.

2.5 Equilibrium

We start by defining the equilibrium.

Definition 1. *Given preferences, endowments, and beliefs, an equilibrium is a collection of allocations* $\left(c_{j,t}, \pi_{i,t}^j\right)$ *and a price system* $\left(r_t, \mu_{E,t}, \{\mu_{R_n,t}\}_{n=1}^N, \mu_{w_E,t}, \mu_{w_\delta,t}, \sigma_{E,t}, \{\sigma_{R_n,t}\}_{n=1}^N\right)$ *for* $j = 1, \ldots, J$ *and such that the processes* $\left(c_{j,t}, \pi_{i,t}^j\right)$ *maximize lifetime utility in Equation* (21) *subject to the dynamic budget condition in* (22) *and all market clear.*

Since the market is complete, we can solve the individual problem using Martingale methods as in Karatzas et al. (1987) and Cox and Huang (1989). The first-order conditions yield

$$c_{j,t} = \left(\kappa_j M_t / \eta_t^j e^{\rho t}\right)^{-1/\gamma}, \tag{23}$$

where κ_j is the Lagrange multiplier from the static optimization problem. The Lagrange multiplier is linked to the initial wealth of agent j. It is convenient to define $\lambda_{j,t} = \eta_t^j / \kappa_j$. By using the optimal consumption in (23) and the market clearing in the commodity market we have the following Proposition.

Proposition 1. *In equilibrium the optimal consumption of agent* $j = 1, \ldots, J$ *is*

$$c_{j,t} = f_{j,t} C_t, \tag{24}$$

where the consumption share, $f_{j,t}$, is

$$f_{j,t} = \frac{\lambda_{j,t}^{\frac{1}{\gamma}}}{\sum_k^J \lambda_{k,t}^{\frac{1}{\gamma}}}. \tag{25}$$

Moreover, the stochastic discount factor, M_t, is

$$M_t = e^{-\rho t} \left(\sum_j^J \lambda_{j,t}^{\frac{1}{\gamma}} \right)^{\gamma} C_t^{-\gamma}. \tag{26}$$

An application of Ito's lemma on the stochastic discount factor, M_t, in (26) and matching the terms with the dynamics in Equation (18) leads to the following Proposition.

Proposition 2. *The equilibrium real rate is*

$$r_t = \rho + \gamma \left(\mu_{C,t} + \sigma_{C,t}' \mathcal{E}_t (\Delta) \right) - \frac{1}{2} \gamma (1 + \gamma) \sigma_{C,t}' \sigma_{C,t} + \frac{1}{2} \left(1 - \frac{1}{\gamma} \right) \mathcal{V}_t (\Delta) \tag{27}$$

and the market prices of risks are

$$\theta_t = \gamma \sigma_{C,t} - \mathcal{E}_t (\Delta) \tag{28}$$

where

$$\mathcal{E}_t (\Delta) = \sum_{j=1}^J f_{j,t} \Delta_t^j \tag{29}$$

is the consumption weighted average and

$$\mathcal{V}_t (\Delta) = \sum_{j=1}^J f_{j,t} \left(\Delta_t^j - \mathcal{E}_t (\Delta) \right)' \left(\Delta_t^j - \mathcal{E}_t (\Delta) \right) \tag{30}$$

is the consumption weighted average total variance of disagreement vector Δ_t^j

The risk-free rate is similar to the standard risk-free rate in a homogeneous beliefs economy with constant relative risk aversion (CRRA) preference with two exceptions. First, the true expected growth rate, $\mu_{C,t}$ is replaced with $\mu_{C,t} + \sigma_{C,t}' \mathcal{E}_t (\Delta)$. This turns out to be equivalent to the consumption share weighed average belief about the consumption growth. This is what Heyerdahl-Larsen and Illeditsch (2020) refer to as the market view. Second, there is an additional term that depends on the consumption share weighted total variance of the disagreement vector, $\mathcal{V}_t (\Delta)$. This term is due to the speculative trade between the agents in the economy and the sign is linked to the value of γ. For $\gamma > 1$, the risk-free rate is higher when there is more disagreement in the economy. This is the channel explored in Ehling et al. (2018). We calculate stock price dynamic using Malliavin calculus. Specifically, it can be shown that the stock price loadings on the shocks are

$$\sigma_{R_i,t} = \theta_t + \frac{E_t\left(\int_t^\infty \mathcal{D}_t\left(M_u D_{i,u}\right) du\right)}{E_t\left(\int_t^\infty M_u D_{i,u} du\right)} \tag{31}$$

where $\mathcal{D}_t x_u = \left(\mathcal{D}_{E,t} x_u, \mathcal{D}_{\delta,t} x_u, \mathcal{D}_{z,t} x_u, \mathcal{D}_{1,t} x_u, \ldots, \mathcal{D}_{N,t} x_u\right)$ denotes the Malliavin derivative of x_u at time t. The next return dynamics is given by the next proposition.

Proposition 3. *The dynamics of stock $i = 1\ldots, N$ is*

$$dR_{n,t} = \mu_{R_n,t} dt + \sigma'_{R_n,t} dw_t \tag{32}$$

where

$$\mu_{R_n,t} = r_t + \theta'_t \sigma_{R_n,t}, \tag{33}$$

and where

$$\sigma_{R_n,t} = \sigma_{D,t} + \theta_t + \frac{E_t\left(\int_t^\infty M_u\left(\mathcal{E}_u\left(\mathcal{D}_t log\left(\eta_u\right)\right) - \gamma \mathcal{D}_t log\left(C_u\right)\right) du\right)}{E_t\left(\int_t^\infty M_u D_{i,u} du\right)}, \tag{34}$$

and where $\mathcal{E}_u\left(\mathcal{D}_t log\left(\eta_u\right)\right)$ is the consumption share weighted average of the Malliavin derivative of the log disagreement process η_u.

As we are interested in the exposure of the agents' wealth to the shocks in the economy, we also need to find the wealth dynamics. The next proposition characterizes the wealth dynamics of the agents in the economy.

Proposition 4. *Let W_t^j be the wealth of agent j at time t with dynamics*

$$dW_t^j = \mu_{W^j,t} W_t^j dt + \sigma'_{W^j,t} W_t^j dw_t, \tag{35}$$

Then the exposures of agent j's wealth to the shocks, $\sigma_{W^j,t}$, are

$$\sigma_{W^j,t} = \theta_t + \frac{E_t\left(\int_t^\infty \mathcal{D}_t\left(M_u C_{j,u}\right) du\right)}{E_t\left(\int_t^\infty M_u C_{j,u} du\right)} \tag{36}$$

From Proposition 4 together with the wealth dynamics in Equation (22) one can calculate the optimal portfolios. As we are interested in the economy wide loading on the factor shock $w_{z,t}$, we instead directly examine the total absolute exposure to the factor shock:

$$\varepsilon_{z,t} = \sum_{j=1}^{J} \mid \sigma_{W^j,z,t} \mid \tag{37}$$

2.6 Model predictions for volatility, correlation and factor exposure

We obtain testable predictions from the model by conducting Monte Carlo simulations aimed at capturing the equilibrium relations between the degree of *factor* disagreement (s_t) and key quantities such as (i) total factor exposure, (ii) the risk of both individual securities and the ETF, and (iii) the average correlation (i.e., diversification benefits) within the ETF.

Each simulation considers an economy with 10 stocks ($N = 10$) and a disagreement parameters $\Delta = 0.2$. Moreover, we assume that each agent has the same Pareto weights, $\frac{1}{\kappa_j}$. This implies that all agents start with the same initial consumption shares. We also assume that the dividend shares of each stock are the same, and that the total dividends are initially 5% of total consumption. We define the ETF as the portfolio of all the 10 stocks. As we consider a frictionless market, there is no intrinsic demand for the ETF in the model. Hence, we make the assumption that agents prefer to trade the ETF instead of the underlying stocks in the ETF if an agent's goal is to take on factor exposure. Underlying this argument is that in a model with a small transaction cost for trading individual stocks, the investors would prefer trading the ETF over the underlying portfolio to get exposure to the factor shock. We average all model-implied results across 400,000 simulated paths of monthly time series that each span 100 years.

Figure 1 plots the the total factor exposure (top-left), the standard deviation of the ETF (top-right), the average stock return correlation (bottom-left) and the average stock return volatility (bottom-right). From the figures, we see that both the exposure and the volatility of the ETF return increases as the disagreement move from stock-specific to factor (high s). As the agents disagree more on the factor risk, larger agents take larger bets on the factor return, and hence total exposure in the economy increases. The higher exposure to the factor risk leads to larger swings in the ETF return as the factor is a common component. This is also reflected by the higher average correlation between the stocks in the ETF (bottom-left plot). The average standard deviation of the stocks in the ETF shows a hump-shaped pattern. The reason for this is that with stock-specific disagreement there are a subset of agents that take large speculative positions on the stock-specific shocks, and this speculative trade increases the volatility. However, this increased volatility "washes" out at the portfolio level. As the economy moves towards more disagreement on the factor, the individual stock return volatility initially drops due to the non-linearity in the model. However, for high levels of s (high factor disagreement), the stock return volatility increases.

Testable predictions. Based on the figures we have the following four model predictions: Increased factor disagreement (higher s_t) leads to:

(a) *Larger flows into the ETF (i.e., more common factor exposure)*;

(b) *Higher ETF-level return volatility*;

(c) *Higher average correlations between the stock returns of securities in the ETF (i.e., lower diversification benefits)*; and

(d) *A weak and non-linear effect on the return volatility of the individual securities in the ETF.*

3 Empirical evidence

This section describes the data and empirical measures used to evaluate the four main predictions of the model outlined in Section 2. Section 3.1 provides an overview of the set of ETFs we use to test the relations between disagreement, fund flows, and risk-neutral volatility and correlation, while Section 3.2 describes our empirical measures of factor and stock-specific disagreement regarding an ETF.

3.1 Data and summary statistics

Our sample begins in January 2012, which is the first month in which ETF Global begins providing granular and comprehensive data on ETF flows and constituents, and ends in December 2020. We employ a set of 14 highly liquid market and sector ETFs as a laboratory to empirically examine the relations between disagreement, flows, and volatility and correlation risk. We focus on ETFs because unlike mutual funds, for whom the disclosures of holdings are relatively infrequent and potentially noisy, granular data on ETF holdings are available at high frequency (i.e., daily). Moreover, since ETFs are exchange-traded securities that are in many cases optioned, it is possible to construct measures of *forward-looking* volatility and correlation risk for an ETF, but similar measures of forward-looking risk are unavailable for mutual funds. Since three of our key predictions are that disagreement impacts a portfolio's volatility and diversification benefits, we exploit ETF options in many of our empirical tests.

As the top panel of Figure 2 shows, ETFs are a relatively nascent security that only began trading in 1993 with the introduction of the SPDR S&P 500 Trust ETF (NYSE: SPY). While ETF trading volumes represented less than 5% of total dollar trading volume in the 1990s, the market has come to represent approximately 25% to 30% of total dollar trading volume since 2010. This rapid increase in the popularity of ETFs represents, in large part, the fact that ETFs provide investors with relatively cheap access to a wide variety of investment factors and "styles." Beyond the fact that granular data on ETF holdings are only available beginning in 2012, there is an additional benefit of starting out sample period at this point in time: our theoretical analyses assumes that the dynamics of disagreement, and consequently flows into and out of an ETF, are stationary. It seems as though ETF trading activity has achieved a stable equilibrium in the time period underlying our analyses.

With these benefits of ETFs in mind, our analysis focuses on a small set of highly liquid US-focused equity ETFs for which options on both the ETF and its constituents are actively traded. As we explain in Section 3.4 and 3.5, focusing on ETFs for which options on both the index and its constituent stocks are actively traded allows us to elicit accurate measures of the *forward-looking* volatility and correlation risk associated with the investment factor, theme, or style the ETF tracks. The 14 ETFs in our sample are SPY (SPDR S&P 500 ETF Trust), DIA (SPDR Dow Jones Industrial Average ETF Trust), QQQ (Invesco QQQ Trust Series 1), XLK (Technology Select Sector SPDR Fund), XLB (Materials Select Sector SPDR), XLE (Energy Select Sector SPDR), XLE (Financial Select Sector SPDR), XLI (Industrial Select Sector SPDR), XLP (Consumer Staples Select Sector SPDR), XLU (Utilities Select Sector SPDR), XLV (Health Care Select Sector SPDR), XLY (Consumer Discretionary Select Sector SPDR), XOP (SPDR S&P Oil & Gas Exploration & Production ETF), and XBI (SPDR S&P Biotech ETF).

Although these 14 ETFs represent only a small *number* of the approximately 2,200 distinct ETFs that are now trading in U.S. markets, the bottom panel of Figure 2 demonstrates that these

14 ETFs represent just under *half* of all dollar trading volume in US ETFs in the recent decade. Moreover, these ETFs represent a variety of investment styles. Three ETFs track broad market index, while 11 track many of the various sectors underlying the U.S. economy. Thus, our sample represent an economically sizable portion of the US equity market.[3]

Table 1 reports a number of summary statistics related to the ETFs that comprise our sample. For instance, the table shows that the largest ETF in our sample is the SPY, which has a net asset value (NAV) of \$227.23b. The smallest ETF is our sample is XOP, the Oil & Gas Exploration & Production ETF that has a net asset value of \$44.90b. Beyond showing relatively large differences in NAVs across ETFs, that table also shows that the market capitalization's of the equities underlying these ETFs also display large differences. For instance, the biotech (healthcare) firms underlying XBI (XLV) have a combined market value of \$710.31b (\$2795.04b).

ETF flows. We measure the net flow into ETF m in month t as

$$\text{Flow}_{m,t} = \frac{\sum_{\tau=1}^{T_t} \text{NetFlow}_{m,t,\tau}}{\sum_{j=1}^{J} \text{ME}_{m,j,t}}. \tag{38}$$

Here, $\text{NetFlow}_{m,t,\tau}$ represents the net flow into ETF m on trading day τ of month t (expressed in dollars and from ETF Global), and T_t captures the total number of trading days in month t. To gain a sense of the economic magnitude of these monthly net flows, and to make flows comparable across ETFs with different market capitalizations, we scale the net flows by the aggregate market capitalization of the J stocks underlying ETF m at the end of month t. The economic intuition behind this choice of scaling is that flows into and out of an ETF should only affect the risks of the ETF and its constituent securities if these flows are large relative to the size of the underlying stocks.

Table 1 reports the average value of the absolute net flows into each ETF over the average month of the sample period. The table shows that approximately \$30b flows into and out the SPY each month, and around one billion dollars moves into and out of XBI and XLB, the biotech and

[3] While a number of other economically large ETFs exist, these ETFs are excluded from our sample because they feature very little option-trading activity at the index level. For instance, while both VOO (Vanguard S&P 500 ETF) and IVV (iShares Core S&P 500 ETF) are two ETFs that have net asset values in excess of \$200b, they are excluded from our sample because there are typically fewer than one thousand options linked to either VOO or IVV traded each day. In contrast, there are typically hundreds of thousands of options linked to SPY traded each day. These differences in options-trading volume at the index level imply that we can elicit risk-neutral moments for SPY, but we are unable to elicit the same moments for either VOO or IVV.

materials ETFs. While the nominal value of these latter flows are an order of magnitude smaller than the flow for the SPY, they represent similar magnitudes relative to the aggregate values of the underlying equities held by each of these three ETFs.

Other summary statistics. Table 1 also reports three additional statistics for each ETF: (i) the proportion of underlying firms that are optioned, or $\mathbb{I}\,[\text{Optioned}]$, (ii) the proportion of underlying firms that are followed by analysts, or $\mathbb{I}\,[\text{Analysts}]$, and (iii) the average number of analysts following each underlying firm, or $\mathbb{E}\,[\text{Analysts}]$. These summary statistics show that the ETFs are well balanced in terms of the degree to which the firms underlying each ETF are optioned and followed by analysts, and in terms of how many analysts follow the average firm in each ETF. This latter fact is useful, as Section 3.2 uses data related to analyst forecasts to construct measures of ETF-level disagreement, which maps to the notion of disagreement in the model (recall equation (9) in Section 2.2).

3.2 Measuring systematic and idiosyncratic disagreement

For our empirical analysis, we need measures of factor and stock-specific disagreement for each ETF. This allows us to empirically determine the extent to which s_t in equation (9) is closer to zero (more stock-specific disagreement) or one (more factor disagreement). Following Buraschi et al. (2014), we estimate these measures using IBES data.

For the *stock-specific* disagreement measure we first calculate the mean absolute value of next period's earnings estimates (Est) for each stock across all analysts. These estimates are subject to the standard filters applied to the unadjusted IBES data file. For instance, we remove all analyst revisions reported after a firm's earnings announcement date. Moreover, we choose each analyst's most recent estimate, removing observations where we know analysts coverage has stopped or where IBES has recommended removal through their own proprietary analysis. Finally, we remove stale information by deleting forecasts that are outstanding for more than 180 days. The measure of stock-specific disagreement surrounding the stocks underlying ETF m at time t, denoted by $\text{StockDisagree}_{m,t}$, is then the weighted sum of the individual security disagreement measures across all stocks in the ETF, or

$$\text{StockDisagree}_{m,t} = \sum_{j=1}^{J} w_{j,m,t} \cdot \left[\frac{\frac{1}{A}\sum_{a=1}^{A} |\text{Est}_{a,j,t} - \overline{\text{Est}}_{j,t}|}{|\overline{\text{Est}}_{j,t}|} \right]. \tag{39}$$

Here, $\text{Est}_{a,j,t}$ is the earnings estimate of analyst a for stock j at time t, $\overline{\text{Est}}_{j,t}$ is the average earnings estimate across all analysts for a given security at time t, and $w_{j,m,t}$ is the weight of stock j in ETF m at time t, drawn from ETF Global. Since these weights are drawn from ETF Global, they typically represent the relative market capitalization of the security in the ETF.

The measure of *factor* disagreement for ETF m at time t, denoted by $\text{FactorDisagree}_{m,t}$, is constructed by first summing the forecasts of all component securities of an index across all analyst employed by a broker. This "bottom up" approach mimics the methodology used by macroeconomic groups at brokers when estimating earnings for the S&P 500 and other indexes (see, e.g., Darrough and Russell (2002)). We then compute the disagreement *across* brokers rather than analysts in the sample,

$$\text{FactorDisagree}_{m,t} = \frac{\sum_{b=1}^{B} |\sum_{j=1}^{J} \text{Est}_{b,j,t} - \overline{\text{Est}}_{m,t}|}{|\overline{\text{Est}}_{m,t}|}, \tag{40}$$

where B is the total number of brokers in the sample, $\text{Est}_{b,j,t}$ is the earnings estimate of broker b for security j, and $\overline{\text{Est}}_{m,t}$ is the equal-weighted average earnings estimate for ETF m across all brokers. Intuitively, this measure captures the extent to which brokers (e.g., Goldman Sachs and J.P. Morgan) disagree about the valuation of an *index* rather than the degree to which analysts disagree about the valuation of an individual constituent of the index. Note that since the weights in our sample are based on the market capitalization of security j versus total market capitalization of all stocks in ETF m, our estimate of $\text{Est}_{b,j,t}$ is simply broker b's earnings-per-share forecast from IBES multiplied by the shares outstanding from COMPUSTAT.

In constructing this measure of factor disagreement, we find that individuals brokers do not necessarily cover all securities in a given ETF. For instance, while almost all brokers cover stocks with large market capitalizations, only some brokers cover stocks with smaller market capitalizations. The dropoff in coverage is not linear, and falls precipitously from over 80% of market capitalization for the top 10 brokers to less than 30% for those outside of the top 10. With these data limitations in mind, we apply the following two filters to the data. First, we restrict our attention to the disagreement among the top 10 brokers in the sample. Second, if a given broker b does not cover a particular stock j at time t, we assume that the broker's estimate for the earnings-per-share of that stock corresponds to the consensus estimate for that security's expected earnings. Since smaller stocks are less covered by analysts, but also have smaller weights in a given ETF, this imputation has a minimal effect on our results.

The top panel of Figure 3 plots the values of stock-specific (or idiosyncratic) and factor disagreement for the SPY, defined following equations (39) and (40). The figure shows that while stock-specific disagreement is often more elevated than systematic disagreement, there are are periods in which most disagreement is systematic in nature. For instance, systematic disagreement regarding SPY exceed stock-specific disagreement in early 2016 in the run up to the 2016 presidential election.

To gauge the *relative* importance of systematic disagreement for a given ETF m in a given month t, we define $S_{m,t}$ as

$$S_{m,t} = \frac{\text{FactorDisagree}_{m,t}}{\text{FactorDisagree}_{m,t} + \text{StockDisagree}}. \tag{41}$$

When this ratio $S_{m,t}$ approaches one, then the majority of the disagreement regarding the prospects of an ETF is related to the prospects of the underlying factor that the ETF tracks. In contrast, when $S_{m,t}$ approach zero, then the majority of the disagreement surrounding an ETF is driven by the (idiosyncratic) prospects of the stocks that constitute the ETF. As such, $S_{m,t}$ provides us with an empirical measure of S_t from equation (9), which governs the amount of disagreement about the dynamics of a stock's idiosyncratic shocks (when $s_t \to 0$) and the amount of disagreement about the common factor (when $s_t \to 1$). We plot the time-series dynamics of $S_{m,t}$ for SPY in the bottom panel of Figure 3. As the top panel of the Figure highlight, the majority of disagreement related to SPY is driven by the idiosyncratic prospects of the constituent stocks of an ETF. However, systematic disagreement exceeds stock-specific disagreement in the period before the 2016 presidential election.[4]

[4]We provide summary statistics for this relative disagreement measure in Table 6 of the Internet Appendix, and report the correlation between these measures of relative disagreement across each pair of ETFs in Table 7 of the Internet Appendix.

3.3 Disagreement and flows

This section shows that increases in disagreement about the common factor underlying each ETF predict an increase in fund flows. That is, when investors face more disagreement about the common component of an investment style rather than stock-specific disagreement about the constituent stocks, then they would rather trade the portfolio than trade the constituent stocks. This supports the first key prediction of the model in Section 2.6, which posits that higher levels of systematic disagreement (i.e., higher values of s_t) predict higher trade flows, as shown in the top-left panel of Figure 1.

We empirically examine the relation between the relative amount of systematic disagreement by estimating the following panel regression

$$\text{Flow}_{m,t} = \gamma_m + \delta_t + \beta_2 \text{S}_{m,t-1} + \beta X_{m,t}^T + \varepsilon_{m,t}. \tag{42}$$

Here, $\text{Flow}_{m,t}$ represents the net flow into ETF m in month t, as defined in equations (), and $\text{S}_{m,t-1}$ corresponds to the relative amount of systematic disagreement regarding ETF m, as defined in equation (41). $X_{m,t}$ represent control variables, and includes one-month lagged flows and ETF returns, as flows in month t may simply arise from investors chasing high returns in month $t-1$. Similarly, this vector of control variables includes the lagged value of the bid-ask spread of the ETF, as investors are less likely to invest in ETFs with larger transactions costs.

We also include ETF and time fixed effects, denoted by δ_t and γ_m, respectively. Time fixed effects absorb common shocks that affect all ETFs simultaneously (e.g., the effect of the Tax Cuts and Jobs Act of 2017, which triggered an inflow of funds into the equity market), while the ETF fixed effects absorb fixed differences in the level of flows and disagreement across ETFs (e.g., the fact flows are unconditionally higher for SPY compared to XBI, the biotechnology sector ETF). Standard errors are clustered by ETF and time and all regressions are estimated at the monthly frequency (see, e.g., Buraschi et al. (2014)). In addition, both our independent and dependent variables are standardized. Recall from equation (3.3) that we scale the net flow into an ETF by the market capitalization of the underlying securities to make the measure comparable across ETFs with different sizes. Given that ETF flows tend to be much larger in magnitude than the market capitalizations of the underlying securities, our flow measures tends to take on small values. Consequently, standardizing the variables allows us to report estimates of the impacts of disagreement on flows that are easier to interpret.

Table 2 reports the results of these panel regressions. Column (1) shows that, without including any ETF and time fixed effects or controls, relatively higher amounts of systematic disagreement are associated with increases in trade flows (i.e., the creation of ETF units). A one-standard-deviation higher amount of systematic disagreement leads to an 0.09 standard deviation higher flow into the ETF next month. Moreover, columns (2) and (3) show that the same results hold true if we control for either ETF or time fixed effects, respectively. Finally, Column (4) shows that controlling for lagged flows, past returns, and trading costs does not alter this result. In particular, while the coefficient on lagged flows is negative (indicating that there is qualitative evidence of mean reversion in flows) and the coefficient on past returns is positively (suggest some of the high flows in month t may be driven by return-chasing investors), neither of these coefficients is statistically significant. There is, however, a negative and significant relation between bid-ask spreads and trading flows. This indicates that net flows are lower in ETFs with larger trading costs.

Combined, these tests validate the intuition developed in our theoretical model. As disagreement increases in in an individual security's factor versus idiosyncratic risk, investors take more concentrated positions in ETFs. This dynamic drives flow into the ETF and correspondingly out of the individual securities, consistent with our model's first prediction.

3.4 Volatility and disagreement

Having shown that disagreement drives flows into ETFs, as predicted by the top-left panel of Figure 1, this section documents that disagreement also drives an increase in ETF-level volatility. This confirms predictions (b) and (d) from Section 2.6, are outlined in the top-right and bottom-right panels of Figure 1, respectively. We establish this relation between disagreement and volatility by first constructing a measure of ETF-level volatility, and then showing that increases in disagreement surrounding the systematic component of ETF earnings drive increases in *forward-looking* ETF-level risk.

Measuring volatility risk. We start by constructing a measure of ETF-level risk using options market data, where forward-looking (risk-neutral) measures of volatility risk are readily available. We follow Bakshi et al. (2003) to construct ETF-level measures of return variation using options data on a given day t. We first express the press of a τ-maturity security that pays the quadratic return as

$$V(t,\tau) = \int_{S(t)}^{\infty} \frac{2\left[1 - \ln(K/S(t))\right]}{K^2} \cdot C\left(t,\tau;K\right) dK + \int_{-\infty}^{S(t)} \frac{2\left[1 + \ln(K/S(t))\right]}{K^2} \cdot P\left(t,\tau;K\right) dK, \quad (43)$$

where K is the strike of either a call $(C\left(t,\tau;K\right))$ or a put $(P\left(t,\tau;K\right))$ option with maturity date $t+\tau$, and $S(t)$ is the current (spot) stock price. Next, we calculate the risk-neutral variance

$$VAR^{Q}(t,\tau) = \exp^{r\tau} V(t,\tau) - \mu_t(\tau)^2. \quad (44)$$

Here, r represents the risk-free rate (drawn form Option Metrics) and $\mu_t(\tau)$ following the definition from Equation 39 of Bakshi et al. (2003). Finally, we convert the daily risk-neutral *variance* into a risk-neutral *volatility*, denoted by $\sigma^{Q}_{m,t,\tau}$ by scaling equation (43) by $\sqrt{252}$.

Equation (44) captures the idea that investors' exposure to a squared return contract is a function of the probability weighted expected return squared across *all* possible share price values. One can back out these values from an infinite string of options in the positive (calls) and negative (puts) return domains (i.e. the volatility surface). Following Buss and Vilkov (2012), we estimate equation (44) for a given ETF over a discretized grid of moneyness (K/S values from 0.33 to 3 by increments of 0.01), using the annual duration volatility surface files from OptionMetrics. Using annual duration options is an important distinction from most other work studying the impact of ETF ownership on intra-ETF correlations and variance (see, e.g., Ben-David et al. (2018) and Da and Shive (2018)) that typically focus on short-maturity measures of risk using realized returns. Long-dated options allow us to estimate changes in expectations of the risk and diversification benefits (i.e., correlation) of an ETF, as well as changes in perceptions of risk. Existing studies characterize the relationship between ETF flow and asset prices as being short-lived and mean-reverting in nature. Consequently, any result we find would be expected to be long-dated due to the extended duration of the options we use.

Figure 4 shows the time series of risk-neutral volatility, estimated using equation (44) for the most prominent ETF in our sample: SPY. In particular, the figure computes the risk-neutral volatility of SPY in two ways. First, equation (44) is estimated using options written on SPY itself (the dashed blue line in the figure). Second, we compute the risk-neutral volatility of each individual firm within SPY, and take the weighted sum of these risk-neutral volatilities across all pairs of firms in the index, where these volatilities are weighted by the importance of a given stock in a given ETF on on a given trading day t (as reported by ETF Global). This procedure results in the dashed red line that is reported in the figure. While these two approaches for calculating the risk-neutral volatility are highly correlated (for instance, each measure of volatility increases surrounding the 2016 presidential election and the COVID-19 crisis), the fact that a wedge exists between the two lines suggests that investors' forward-looking perceptions of the intra-ETF correlation is both less than one (i.e., investors do not believe the stocks within the S&P 500 are perfectly positively correlated) and varies over time.

Following the intuition underlying our economic model in Section 2, we implement a regression analysis to examine whether disagreement regarding the relative contribution of the systematic risk of stock returns also drives (part of) the cross-sectional differences in the forward-looking volatility of an ETF in a given month. Section 3.5 then explores whether disagreement is also associated with the intra-ETF correlation between stock, as our model also predicts.

Regression analysis. We examine the relation between the relative importance of systematic disagreement and ETF-level volatility risk by estimating the following panel regression

$$\sigma^Q_{m,t,\tau} = \gamma_m + \delta_t + \beta_2 S_{m,t-1} + \beta X^T_{m,t} + \varepsilon_{m,t}. \tag{45}$$

As in regression (42),$\mathrm{S}_{m,t-1}$ corresponds to the relative amount of systematic disagreement regarding ETF m, as defined in equation (41). When estimating this equation, we measure the risk-neutral volatility of ETF m at time t by focusing on options with $\tau = 30$ days to maturity. The controls we use are the exact same as in regression (42), and include lagged flows, and ETF returns and bid-ask spreads. We also include ETF and time fixed effects, denoted by δ_t and γ_m, respectively, and standardized variables so that each point estimate can be interpreted as the effect of a one-standard deviation change in the variable of interest. The coefficient of interest is that on the relative disagreement measure. Similar to table 2, we expect β_2 to be positive, as increases in factor disagreement would lead to a higher variance of the ETF overall.

Table 3 reports the results of these panel regressions. Columns (1) shows that increases in the systematic portion of disagreement result in higher ETF-level volatility, as predicted by our model. Column (1) establishes this fact without including any fixed effects or controls in the regression. Columns (2) and (3) add ETF and ETF and date fixed effects, respectively, and show that higher systematic disagreements continues to suggest increases in forward-looking volatility. Finally, Column (4) adds a host of control variables to the specification from Column (3), and once again documents that increases in factor disagreement are associated with strong increases in volatility. A one standard deviation higher factor disagreement measure leads to a 0.08 standard-deviation higher volatility, even when we control for lagged ETF flows, returns, and trading costs. These tests validate the intuition developed in our theoretical model: as disagreement in the factor increases across agents, volatility increases by a statistically and economically large amount.

3.5 Intra-ETF correlation risk

In our model the volatility of individual security cash flows are constant. The relationship between flow into the factor instrument and factor volatility is entirely driven by changing correlations between the stocks composing the factor. This implies that there should be a strong connection between flows induced by greater factor disagreement and correlations. In the context of our risk-neutral measure of implied volatility, one would therefore assume that correlation *risk* is related to flow. In this section we empirical test, and confirm, this prediction.

Measuring correlation risk. To construct a measure of the *forward-looking* diversification benefits of holding an ETF, we start by considering the definition of portfolio variance for ETF m on day t using risk-neutral measures of volatility constructed using options with τ days to maturity

$$\left(\sigma^Q_{m,t,\tau}\right)^2 = \sum_{i=1}^{N}\sum_{j=1}^{N} w_{i,m,t} w_{j,m,t} \sigma^Q_{i,t,\tau} \sigma^Q_{j,t,\tau} \rho^Q_{i,j,t,\tau}. \tag{46}$$

Here, $\sigma^Q_{m,t,\tau}$ denotes the risk-neutral volatility of ETF m, $\sigma^Q_{i,t,\tau}$ is the risk-neutral volatility of security i, $w_{i,m,t}$ denotes the weight of security i in ETF m, obtained from ETF Global, and $\rho^Q_{i,j,t,\tau}$ represents the risk-neutral correlation between stocks i and j on day t with τ days to maturity. We obtain the risk-neutral volatility of both the index and the constituent stocks via equation (44).

While exchange-traded options prices provide us with readily observable estimates of $\sigma^Q_{m,t,\tau}$ and $\sigma^Q_{i,t,\tau}$, the options market does not provide us with exchange-traded claims that deliver the correlation between a pair of securities at maturity. Consequently, we obtain the *average* risk-neutral correlation the between pairs of securities that constitute ETF m by "inverting" equation (46) for the average value of $\rho^Q_{i,j,t,\tau}$. That is, we define

$$\overline{\rho}^Q_{m,t,\tau} = \frac{\left(\sigma^Q_{m,t,\tau}\right)^2}{\sum_{i=1}^N \sum_{j=1}^N w_{i,m,t} w_{j,m,t} \sigma^Q_{i,t,\tau} \sigma^Q_{j,t,\tau}}. \tag{47}$$

as the risk-neutral correlation of stocks contained in ETF m at time t. This measure essentially reflects the time-varying wedge between the measures of index-level volatility computed using index-level and individual stock options in Figure 4.

Figure 5 displays the monthly time-series variation for the average 30-day risk-neutral correlations ($\rho^Q_{m,t,\tau=30}$) underlying two prominent ETFs in the sample: SPY and XLF, an ETF that tracks financial stocks. There are four key takeaways from this figure. First, the estimated values of $\overline{\rho}^Q_{m,t,\tau=30}$ clearly satisfy the requirement that $\rho^Q_{m,t,\tau=30} \in [0, 1]$. Second, figure shows that these risk-neutral correlations vary substantially over time. While the average correlation between S&P 500 stocks was between 0.50 and 0.70 in 2012 (depending on the maturity of the options used to compute $\overline{\rho}^Q_{m,t,\tau=30}$ in equation (47), these correlations dropped a low as 0.20 during the economic expansion of the 2010s before rising back to around 0.70 during the onset of the COVID-19 crisis. Third, there is also a large degree of *cross-sectional* heterogeneity between the risk-neutral correlations across ETFs at a given point in time. For instance, while stocks in the S&P 500 shared an average risk-neutral correlation of about 0.50 to 0.70 in 2012, financials had an average risk neutral correlation of between 0.70 to 0.90. Thus, ETF-level correlation risk can vary both over time and between different investment styles.

The final takeaway from Figure 5 is that there are stark differences in the *term structure* of risk-neutral correlations. Notably, while short-term measures of intra-ETF correlations (estimated using 30-day options) tend to be "fast-moving" and vary substantially over time, long-term measures of intra-ETF correlations (estimated using 365-day options) are relatively "slow-moving" and consequently vary less over time. However, these differences in the levels of correlations tend to diverge during good times and compress during bad times, such as the economic and financial shock induced by COVID-19. We turn to the implications of disagreement for the term structure of forward-looking correlations in Section 3.5.

Regression analysis. To begin, we examine how disagreement regarding the systematic component of ETF returns may drive variation in the *level* of intra-ETF correlation risk by estimating the following panel regression

$$\overline{\rho}^{Q}_{m,t,\tau} = \delta_t + \gamma_m + \beta_2 S_{m,t-1} + \beta X^{T}_{m,t-1} + \varepsilon_{m,t}. \tag{48}$$

Here, $\overline{\rho}^{Q}_{m,t,\tau}$ represents the average risk-neutral correlation for ETF m at time t, estimated using 30-day options, and $S_{m,t-1}$ is the amount of factor-level disagreement related to ETF m at time t, as defined in equation (41), and the vector $X_{m,t-1}$ controls for ETF flows, returns, and bid-ask spreads. The regression includes both time fixed effects (δ_t) that account for common shocks that impact all ETFs at a given point in time, and ETF fixed effects (γ_m) that account for unconditional differences in correlation risk across ETFs (e.g., the difference in the level of correlation between SPY and XLF in Figure 5. Finally, and similar to the tables above, both the independent and dependent variables are standardized for proper interpretation.

The results of estimating equation (48) are reported in Table 4. The table highlights three key takeaways. First, as shown in Column (1) there is a strong predictive relationship between correlation and lagged disagreement. A one-standard-deviation increase in relative factor disagreement results in a 0.25 standard-deviation increase in $\overline{\rho}^{Q}_{m,t,30}$, which empirically supports the intuition underlying model prediction (c) in Section 2.6. Specifically, these empirical results support the prediction illustrated by the bottom-left panel of Figure 1. Through the lens of the model, this association arises because more concentrated trading in the ETF (i.e., a greater exposure to the systematic risk factor) increases the covariation between related securities, and makes the market more susceptible to shocks that impact this systematic factor. In Columns (2) and (3) we add ETF and month fixed effects, respectively, to the regression; while the coefficient are still statistically significant the economic magnitude decreases with the addition of ETF fixed effects. Now a one-standard-deviation increase in the relative factor disagreement predicts a 0.06 standard-deviation increase in $\overline{\rho}^{Q}_{m,t,30}$. This result indicates that the relationship holds true both dynamically and in the cross-section. In columns (4) through (6) we add lagged ETF flows, returns and bid-ask spreads as controls. The statistical and economic significance of the relationship between lagged disagreement and correlation remains. Collectively, the evidence in this section further validates the economic mechanisms of our model.

Correlation term-structure. One of the underlying assumptions of the model is the relative share of factor disagreement is mean reverting. One could directly test the mean-reverting properties of $S_{m,t}$ itself. However, due to its relatively short time-series and high persistence, the analysis may be unreliable. Alternatively, one could look at how long the market anticipates the higher correlation to persist using correlations implied by options of different horizons. We therefore test how shifts in disagreement propagate through the term structure of forward-looking correlations by running the following panel regression

$$\overline{\rho}^{Q}_{m,t,\tau_2} - \overline{\rho}^{Q}_{m,t,\tau_1} = \delta_t + \beta_2 S_{m,t-1} + \beta X^{T}_{m,t-1} + \varepsilon_{m,t}. \tag{49}$$

Here $\overline{\rho}^{Q}_{m,t,\tau_2}$ and $\overline{\rho}^{Q}_{m,t,\tau_1}$ represent average risk-neutral correlations for ETF m at time t, estimated using options of two different horizons, τ_2 and τ_1, respectively, with $\tau_2 > \tau_1$. This difference in risk-neutral correlations is regressed onto the factor-level disagreement related to ETF m at time t as defined in equation (41), and the vector $X_{m,t-1}$ controls for ETF flows, returns, and bid-ask spreads. Given that the dependent variable is already a measured as a difference, we add time (monthly) fixed-effects to the regression.

The results of estimating equation (49) are reported in Table 5. The columns highlight the correlation spreads that are used as independent variables for each regression. Furthermore, we do not standardize the independent variable in order to highlight the economic significance of shifts in the term structure. First, as highlighted in Columns (1) and (2), most of the activity takes place in the front end of the correlation term structure, between horizons of 30-days and 6-months. This contrasts with the duration of correlation and volatility effects that result from the arbitrage-based mechanisms of, e.g., Ben-David et al. (2018) and Da and Shive (2018), which tend to persist less than 3 months.

To better understand the magnitude of these effects, we have also listed the means of each spread at the bottom of Table 5. On average the correlation spread is steepest between 91 and 30-days (i.e., Column (1)), with implied correlation at 91 days being almost 0.12 correlation points higher, and relatively flat for longer horizons. Relative to the baseline spreads, however, the effects from a one-standard-deviation increase in relative disagreement is most pronounced on the spread between 182 and 91-days. This strongly suggests that systematic disagreement and the resulting flows into the ETF drive longer-dated anticipated correlation. Finally relative disagreement has little predictability on the correlation term structure spread at longer horizons. The expectations hypothesis suggests that the market therefore anticipates a mean reversion of any shock to the term structure induced by a shift in relative disagreement. Given the relationships we have found between lagged lagged $S_{m,t}$, and volatility and correlation risk, this implies a mean reverting $S_{m,t}$ process.

4 Conclusion

This paper studies how factor and stock-specific disagreement affect asset prices and risk. We start by building a theoretical model to examine the implications of belief polarization, focusing on the interplay between heterogeneity in subjective expectations about the stock-specific and common components of expected returns. We consider a pure exchange economy with multiple Lucas trees that are exposed to both factor risk and stock-specific shocks, and we allow agents to disagree about both dimensions of returns. As such, our model features periods of strong disagreement about the factor, and periods when disagreement about stocks dominates.

Our model predicts that (i) factor disagreement increases the exposure of investors in the economy to the assets that are most closely aligned with the systematic risk factor; (ii) factor disagreement increases the return volatility of financial instruments aligned with the common factor (e.g., ETFs); (iii) this increase in volatility is caused by the increased correlations between stocks that compose the factor; and (iv) there is a U-shaped relationship between factor disagreement and the volatility of the individual securities composing the factor.

We then use the return and flow dynamics of ETFs and their underlying securities to test these hypotheses. In keeping with the model's predictions we first find that disagreement is strongly related to ETF flows: when the proportion of factor-to-stock-specific disagreement is one-standard-deviation higher, there is a 0.08 standard-deviation higher flow into the ETF. Second, disagreement also relates closely to forward-looking and long-dated volatility: a one-standard-deviation increase in relative factor disagreement measure leads to volatility rising by 0.06 standard deviations. Finally, we find that increased factor disagreement is also strongly related to a higher long-run correlations. Specifically, an increase in the relative magnitude of factor disagreement *reduces* the diversification benefits of holding the stocks underlying the ETF. Taken together, our theoretical and empirical result highlight the first-order effect of disagreement on trade flows, and show how flows impact cross-sectional differences in volatility and correlation risk.

References

Ajinkya, B. B., Atiase, R. K., Gift, M. J., 1991. Volume of trading and the dispersion in financial analysts' earnings forecasts. Accounting Review pp. 389–401.

Atmaz, A., Basak, S., 2018. Belief dispersion in the stock market. The Journal of Finance 73, 1225–1279.

Bakshi, G., Kapadia, N., Madan, D., 2003. Stock return characteristics, skew laws, and the differential pricing of individual equity options. The Review of Financial Studies 16, 101–143.

Baltussen, G., van Bekkum, S., Da, Z., 2019. Indexing and stock market serial dependence around the world. Journal of Financial Economics 132, 26–48.

Barberis, N., Shleifer, A., Wurgler, J., 2005. Comovement. Journal of financial economics 75, 283–317.

Basak, S., 2000. A model of dynamic equilibrium asset pricing with heterogeneous beliefs and extraneous risk. Journal of Economic Dynamics and Control 24, 63–95.

Basak, S., 2005. Asset pricing with heterogeneous beliefs. Journal of Banking and Finance 29, 2849–2881.

Ben-David, I., Franzoni, F., Moussawi, R., 2012. Hedge fund stock trading in the financial crisis of 2007–2009. The Review of Financial Studies 25, 1–54.

Ben-David, I., Franzoni, F., Moussawi, R., 2018. Do etfs increase volatility? The Journal of Finance 73, 2471–2535.

Berkman, H., Dimitrov, V., Jain, P. C., Koch, P. D., Tice, S., 2009. Sell on the news: Differences of opinion, short-sales constraints, and returns around earnings announcements. Journal of Financial Economics 92, 376–399.

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

Bessembinder, H., Chan, K., Seguin, P. J., 1996. An empirical examination of information, differences of opinion, and trading activity. Journal of Financial Economics 40, 105–134.

Bhamra, H. S., Uppal, R., 2014. Asset prices with heterogeneity in preferences and beliefs. Review of Financial Studies 27, 519–580.

Bond, P., Garcia, D., 2021. The equilibrium consequences of indexing. Forthcoming, RFS .

Brogaard, J., Ringgenberg, M. C., Sovich, D., 2019. The economic impact of index investing. The Review of Financial Studies 32, 3461–3499.

Brown, G. W., Howard, P., Lundblad, C. T., 2021. Crowded trades and tail risk. Available at SSRN 3326802 .

Buraschi, A., Jiltsov, A., 2006. Model uncertainty and option markets with heterogeneous beliefs. Journal of Finance 61, 2841–2897.

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

Buss, A., Vilkov, G., 2012. Measuring equity risk with option-implied correlations. The Review of Financial Studies 25, 3113–3140.

Chen, J., Hong, H., Stein, J. C., 2002. Breadth of ownership and stock returns. Journal of financial Economics 66, 171–205.

Cox, J., Huang, C., 1989. Optimal Consumption and Portfolio Policies when Asset Prices Follow a Diffusion Process. Journal of Economic Theory 83, 33–83.

Cujean, J., Hasler, M., 2017. Why does return predictability concentrate in bad times? Journal of Finance 72, 2717–2758.

Cvitanic, J., Jouini, E., Malamud, S., Napp, C., 2012. Financial markets equilibrium with heterogeneous agents 16, 285–321.

Cvitanic, J., Malamud, S., 2011. Price impact and portfolio impact. Journal of Financial Economics 100, 201–225.

Da, Z., Shive, S., 2018. Exchange traded funds and asset return correlations. European Financial Management 24, 136–168.

Daniel, K., Klos, A., RottkeD, S., 2021. The dynamics of disagreement .

Darrough, M. N., Russell, T., 2002. A positive model of earnings forecasts: Top down versus bottom up. The Journal of Business 75, 127–152.

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

Detemple, J. B., Murthy, S., 1994. Intertemporal asset pricing with heterogeneous beliefs. Journal of Economic Theory 62, 294–320.

Diether, K. B., Malloy, C. J., Scherbina, A., 2002. Differences of opinion and the cross section of stock returns. The Journal of Finance 57, 2113–2141.

Dumas, B., Kurshev, A., Uppal, R., 2009. Equilibrium portfolio strategies in the presence of sentiment risk and excess volatility. Journal of Finance 64, 579–629.

Ehling, P., Gallmeyer, M., Heyerdahl-Larsen, C., Illeditsch, P., 2018. Disagreement about inflation and the yield curve. Journal of Financial Economics 127 (3), 459 – 484.

Ellul, A., Jotikasthira, C., Lundblad, C. T., 2011. Regulatory pressure and fire sales in the corporate bond market. Journal of Financial Economics 101, 596–620.

Goetzmann, W. N., Massa, M., 2005. Dispersion of opinion and stock returns. Journal of Financial Markets 8, 324–349.

Harrison, J., Kreps, D., 1978. Speculative investor behavior in a stock market with heterogeneous expectations. Quarterly Journal of Economics 92, 323–336.

Heyerdahl-Larsen, C., Illeditsch, P., 2020. The market view. Working paper .

Huang, S., O'Hara, M., Zhong, Z., 2021. Innovation and informed trading: Evidence from industry etfs. The Review of Financial Studies 34, 1280–1316.

Jouini, E., Napp, C., 2007. Consensus consumer and intertemporal asset pricing with heterogeneous beliefs. Review of Economic Studies 74, 1149 – 1174.

Karatzas, I., Lehoczky, J., Shreve, S., 1987. Optimal portfolio and consumption decisions for a "small investor" on a finite horizon. SIAM journal on control and optimization 25, 1557–1587.

Park, C., 2005. Stock return predictability and the dispersion in earnings forecasts. The Journal of Business 78, 2351–2376.

Scheinkman, J. A., Xiong, W., 2003. Overconfidence and speculative bubbles. Journal of Political Economy 111, 1183–1219.

Xiong, W., Yan, H., 2010. Heterogeneous expectations and bond markets. Review of Financial Studies 23, 1433–1466.

Zapatero, F., 1998. Effects of financial innovations on market volatility when beliefs are heterogeneous. Journal of Economic Dynamics and Control 22, 597–626.

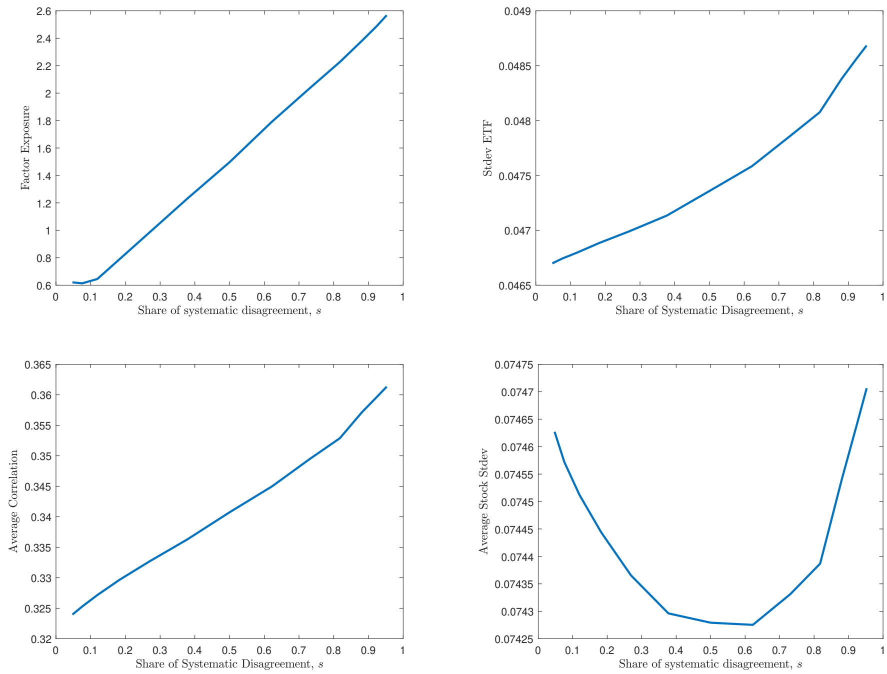

Figure 1: The figures shows the factor exposure (top-left), the standard deviation of the ETF (top-right), the average stock return correlation (bottom-left) and the average stock return standard deviation (bottom-right) as a function of the share of systematic disagreement s. The figures are based on a 10 stock economy. All agents have the same consumption shares and each dividend share are the same for each stock. The total dividend share is 5%. The figures is based on the following parameters: $\gamma = 2$, $\rho = 0.02$, $\mu_E = 0.02$, $\sigma_E = 0.03$, $\sigma_Z = 0.03$, $\mu_Z = -0.5\sigma_Z^2$, $\sigma_n = 0.06$, $\mu_n = -0.5\sigma_n^2$ $\Delta = 0.02$, $\kappa_\delta = 0.1$, $\sigma_\delta = 0.2$, $\bar{\delta} = 0$. The results are generated by Monte-Carlo simulations based on 400,000 paths of monthly observations of length of 100 years.

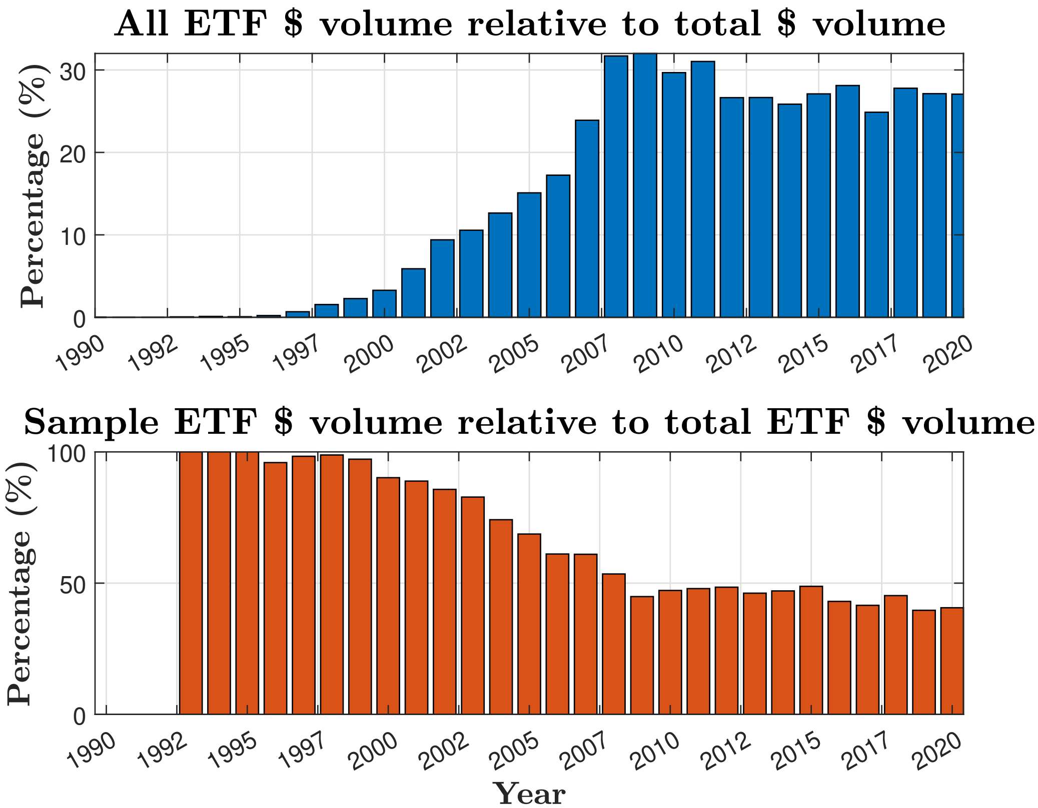

Figure 2: The figure displays the percentage of dollar-trading volume in ETFs relative to dollar-trading volume for the U.S. equity market for the period ranging from 1900 to 2020. ETFs are identified as securities in the CRSP Monthly dataset that have a share code (SHRCD) of 73. We then compute the monthly dollar trading volume of (i) all ETFs, and (ii) all securities in the CRSP Monthly universe, and aggregate these monthly trading volumes to the annual frequency within each trading year. The top panel of the figure then reports the percentage of ETF-related dollar trading volume, summed across all U.S. ETFs, relative to the aggregate amount of dollar-trading volume across all U.S. securities. The bottom panel of the figure reports the percentage of ETF-related dollar trading volume, summed across the 14 ETFs in our sample, relative to the aggregate amount of dollar-trading volume across all U.S. ETFs.

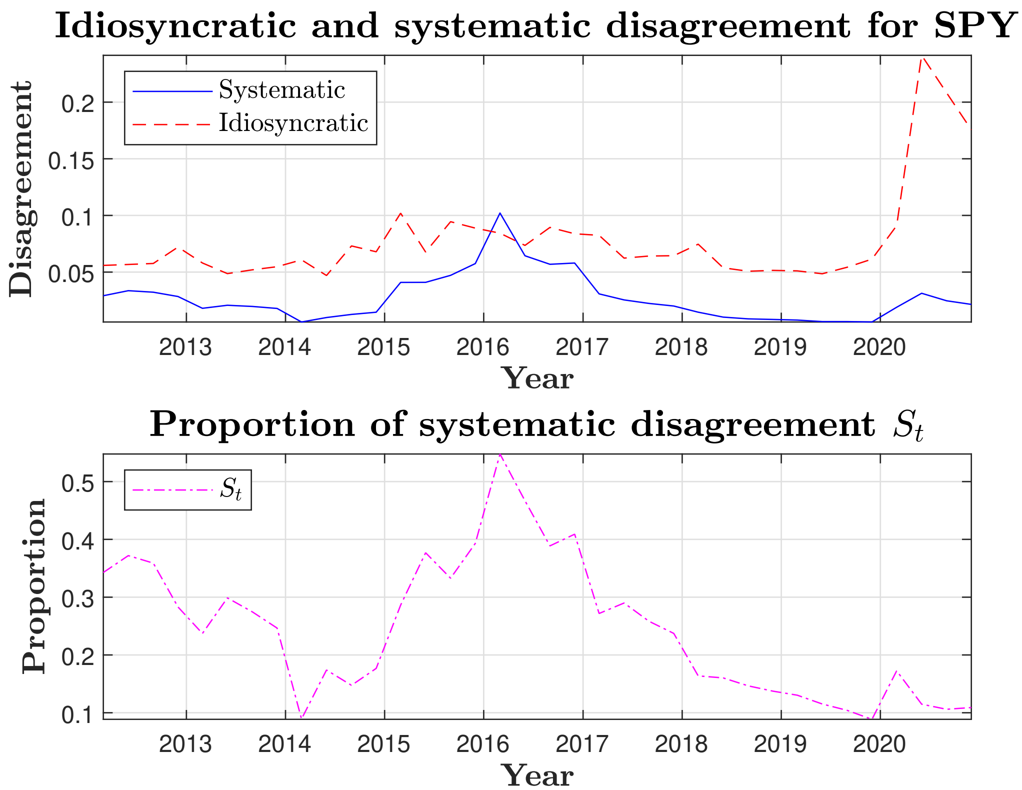

Figure 3: The top panel of the figure reports the levels of systematic (factor) and idiosyncratic disagreement for the SPY – an ETF that tracks the S&P 500 – as measured using equations (39) and (40) respectively. The bottom plots the proportion of systematic disagreement relative to total disagreement, as defined by equation (41). For the purpose of visualizing the data, we aggregate each measure of disagreement to the quarterly frequency by computing the mean value of each disagreement in a given quarter. The sample period ranges from 2012 to 2020.

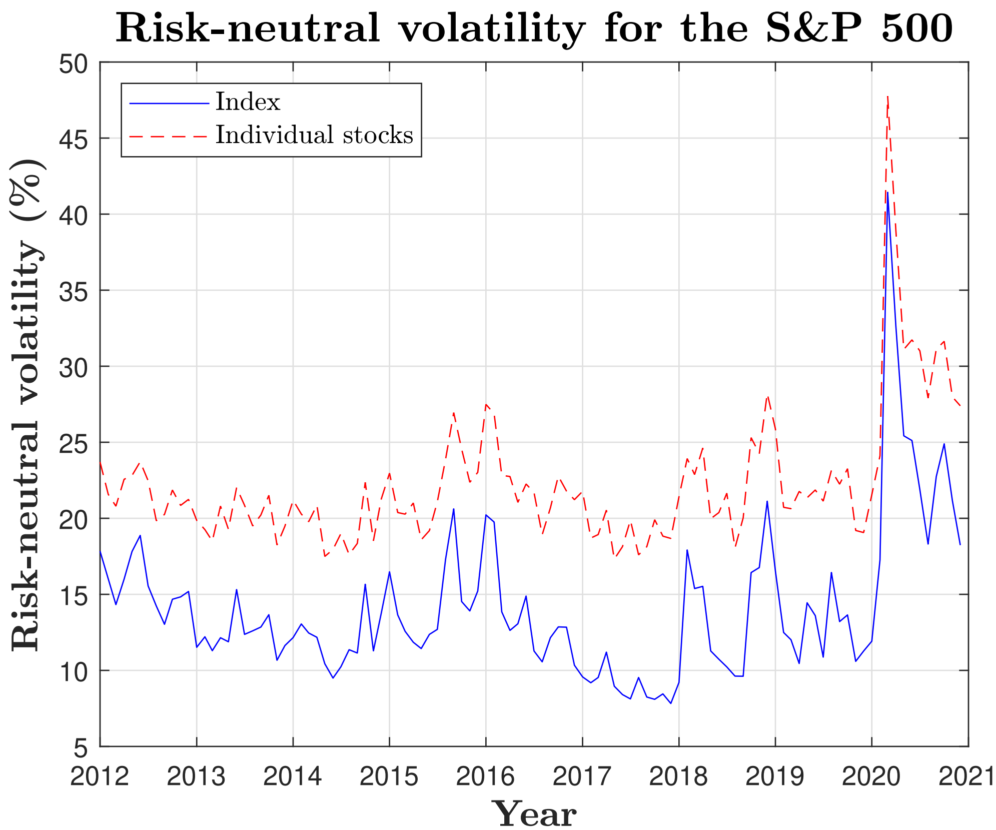

Figure 4: The figure displays the risk-neutral volatility for the SPY, which tracks the S&P 500 stock market index. The risk-neutral volatility is computed in two ways. First, the solid blue line reports the volatility obtained by estimating equation (44) using ETF-linked options. Second, the dashed red line reports the volatility obtained by estimating equation (44) using individual stock options, for each stock i and j in ETF m, and then computing the index-level volatility as $\sum_{i=1}^{N}\sum_{i=1}^{N} w_{i,m,t}w_{j,m,t}\sigma^{Q}_{i,t,\tau}\sigma^{Q}_{j,t,\tau}$, where $w_{i,m,t}$ and $w_{j,m,t}$ is the weight of firm i or j in ETF m at time t, respectively, and each risk-neutral volatility (i.e., $\sigma^{Q}_{i,t,\tau}$) is obtained by taking the square root of the risk-neutral variance in equation (44). These risk-neutral volatilities are calculated on a daily basis, scaled to represent annualized volatilities, and aggregated to the monthly frequency by computing the average risk-neutral volatility within each month. In all the calculations above, we set $\tau = 30$ such that we estimate each risk-neutral volatility with 30-day options. Finally, the sample period spans January 2012 and December 2020.

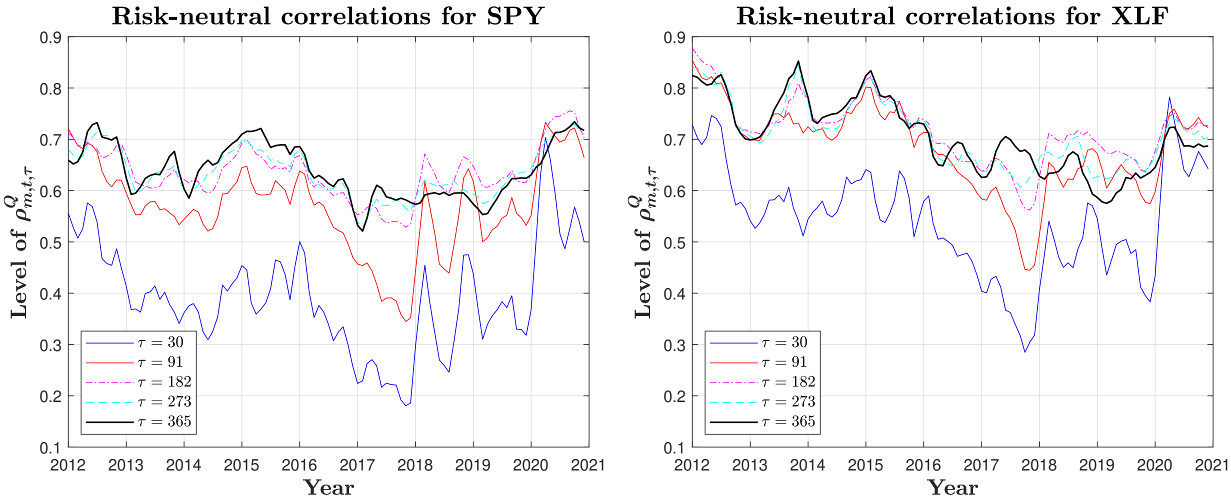

Figure 5: The figure displays the average risk-neutral correlation for the SPY, which tracks the S&P 500 stock market index, the XLF ETF, which tracks the returns of financial stocks. The figure reports the τ = 30-, 91-, 182-, 273-, and 365-day-ahead risk-neutral correlation of each index, obtained by solving equation (47) to obtain $\overline{\rho}^{Q}_{m,t,\tau}$ for each ETF m in each month t. For the purpose of visualizing the resulting risk-neutral correlations, we apply a moving-average filter to each monthly time-series of correlations, and report the average correlation over a window of $[-1, 1]$ months around each month t. The sample period spans January 2012 to December 2020.

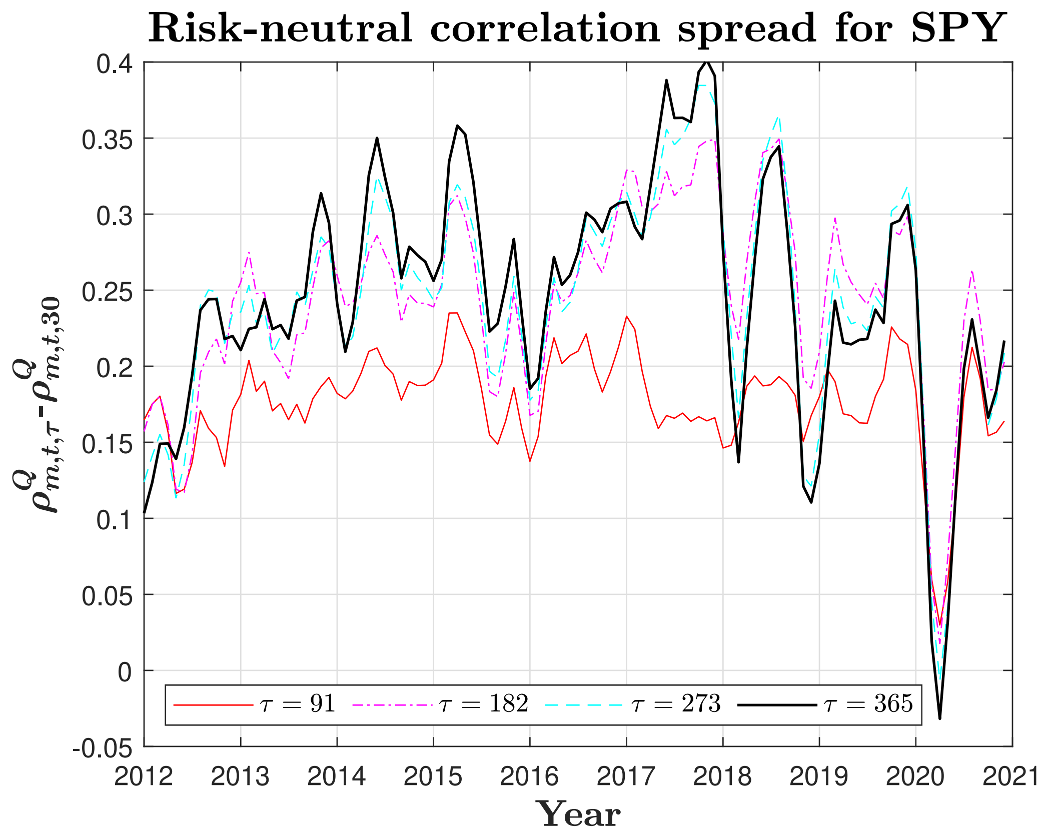

Figure 6: The figure reports the time-series of the spread in the average risk-neutral correlations among the constituents of the SPY ETF, an ETF that tracks the returns of the S&P 500 index. Specifically, the risk-neutral correlation $\overline{\rho}^Q_{m,t,\tau}$ represented by equation (47) is estimated using options with maturities of $\tau \in \{30, 91, 182, 273, 365\}$ days to maturity. With these risk-neutral correlations in hand, the figure then reports $\overline{\rho}^Q_{m,t,\tau} - \overline{\rho}^Q_{m,t,30}$ for $\tau \in \{91, 182, 273, 365\}$. As such, the figure displays the difference between various measures of the intra-ETF "long-run" correlations and the market's perception of the intra-ETF "short-run" correlation. For the purpose of visualizing the resulting risk-neutral correlation spreads, we apply a moving-average filter to each monthly time-series of correlation spreads, and report the average correlation spread over a window of $[-1, 1]$ months around each month t. The sample period spans January 2012 to December 2020.

Table 1: The table reports summary statistics for the 14 ETFs in our sample. It reports the ticker of each ETF, alongside the benchmark that each ETF tracks (denoted by "Style"). The net asset value of each fund is represented by "NAV ($b)", and is reported by ETF Global. Similarly, "ME ($b)" represents the average total market value represented by the stocks underlying each ETF over the sample period, and is constructed from CRSP Monthly data. "|Flow|" reports the average amount of (absolute) flow into and out of each ETF, on average, over the sample period, and is also measured in billions of dollars from ETF Global. $I\,[\text{Optioned}]$ and $I\,[\text{Analysts}]$ represent the proportion of stocks underlying each ETF that are optioned and followed by analysts, respectively, which $E\,[\text{Analysts}]$ reports the average number of analysts following each firm in a given ETF. The sample period ranges from 2012 to 2020.

ETF	Style	NAV ($b)	ME ($b)	|Flow| ($b)	$I\,[\text{Optioned}]$	$I\,[\text{Analyst}]$	$E\,[\text{Analyst}]$
SPY	Market	227.23	20227.72	29.36	1.00	0.99	16.27
DIA	Dow Jones	200.56	5840.02	2.30	1.00	0.99	21.36
QQQ	Nasdaq	133.55	6219.90	7.01	1.00	1.00	19.58
XLK	Technology	55.02	4760.55	1.61	1.00	1.00	21.22
XLB	Materials	49.77	577.13	0.91	1.00	1.00	15.18
XLE	Energy	69.13	1346.49	2.17	1.00	0.99	24.22
XLF	Financials	23.29	2803.44	3.14	1.00	1.00	16.06
XLI	Industrials	59.98	1979.36	2.02	1.00	1.00	15.12
XLP	Cons. Staples	50.01	1942.78	1.82	1.00	0.99	14.53
XLU	Utilities	48.29	625.77	1.86	1.00	1.00	10.61
XLV	Health Care	72.35	2795.04	1.90	1.00	1.00	15.58
XLY	Cons. Disc.	85.73	2408.32	1.60	1.00	0.99	19.26
XOP	Oil & Gas	44.90	1163.23	1.17	1.00	0.99	17.04
XBI	Biotech	104.43	710.31	0.96	1.00	0.99	7.33

Table 2: The table documents how changes in relative factor disagreement ($S_{m,t}$) about an ETF are associated with trade flows into and out of the ETF. These results are obtained by estimating the panel regression outlined in equation (42), the net flows into an ETF are constructed according to equation (3.3), and the measure of relative disagreement are obtained via equation (41). Each regression also controls for one-month lagged trade flows, ETF-level returns and bid-ask spreads, as well as combinations of ETF and time (i.e., month) fixed effects. The sample period underlying this regression ranges from 2012 to 2020.

$S_{m,t-1}$	0.0933***	0.0845***	0.0860**	0.0740**	0.0808**	0.0842***
	[3.06]	[4.07]	[2.31]	[2.47]	[2.89]	[3.03]
$Flow_{m,t-1}$				−0.0576	−0.0688	−0.0652
				[−1.37]	[−1.52]	[−1.46]
$Ret_{m,t-1}$					0.1000	0.0982
					[1.20]	[1.17]
Bid-Ask$_{m,t-1}$						−0.0497***
						[−3.50]
Date FE	No	No	Yes	Yes	Yes	Yes
ETF FE	No	Yes	No	Yes	Yes	Yes
Observations	1,362	1,362	1,362	1,362	1,362	1,362
R^2	0.0089	0.0164	0.1217	0.1324	0.1372	0.1390

Table 3: The table documents how changes in relative factor disagreement ($S_{m,t}$) about an ETF are associated with the *forward-looking* risk of an ETF, measured using option-implied volatility. These results are obtained by estimating the panel regression outlined in equation (45), the forward-looking risk of an ETF is obtained via equation (44), and the measure of relative disagreement are obtained via equation (41). Each regression also controls for one-month lagged trade flows, ETF-level returns and bid-ask spreads, as well as combinations of ETF and time (i.e., month) fixed effects. The sample period underlying this regression ranges from 2012 to 2020.

$S_{m,t-1}$	0.1329^{*}	0.1006^{*}	0.0741^{**}	0.0736^{**}	0.0690^{**}	0.0624^{**}
	[2.07]	[2.09]	[2.49]	[2.47]	[2.36]	[2.42]
$\text{Flow}_{m,t-1}$				0.0177	0.0252	0.0182
				[0.95]	[1.24]	[1.08]
$\text{Ret}_{m,t-1}$					-0.0670	-0.0635
					[-1.32]	[-1.31]
$\text{Bid-Ask}_{m,t-1}$						0.0972^{***}
						[8.84]
Date FE	No	No	Yes	Yes	Yes	Yes
ETF FE	No	Yes	Yes	Yes	Yes	Yes
Observations	1,362	1,362	1,362	1,362	1,362	1,362
R^2	0.0169	0.5090	0.8924	0.8927	0.8947	0.9010

Table 4: The table documents how changes in relative factor disagreement ($S_{m,t}$) about an ETF are associated with the *forward-looking* correlation of an ETF, measured using equation (47). The results are obtained by estimating the panel regression outlined in equation (48), and the measure of relative disagreement are obtained via equation (41). Each regression also controls for one-month lagged trade flows, ETF-level returns and bid-ask spreads, as well as combinations of ETF and time (i.e., month) fixed effects. The sample period underlying this regression ranges from 2012 to 2020.

$S_{m,t-1}$	0.2565***	0.0683*	0.2594***	0.0418***	0.0383***	0.0314**
	[5.84]	[1.91]	[5.36]	[2.84]	[2.62]	[2.17]
$Flow_{m,t-1}$				0.0128	0.0185	0.0112
				[0.86]	[1.22]	[0.75]
$Ret_{m,t-1}$					−0.0511**	−0.0474**
					[−2.35]	[−2.25]
$\text{Bid-Ask}_{m,t-1}$						0.1017***
						[3.44]
Date FE	No	No	Yes	Yes	Yes	Yes
ETF FE	No	Yes	No	Yes	Yes	Yes
Observations	1,362	1,362	1,362	1,362	1,362	1,362
R^2	0.0641	0.5104	0.3648	0.8109	0.8121	0.8192

Table 5: The table documents how changes in relative factor disagreement ($S_{m,t}$) about an ETF are associated with the term-structure of correlation of an ETF, measured by subtracting option implied *forward-looking* correlation at different horizons. The horizons used are listed in the header of the tables. The results are obtained by estimating the panel regression outlined in equation (49), and the measure of relative disagreement are obtained via equation (41). Each regression also controls for one-month lagged trade flows, ETF-level returns and bid-ask spreads, as well as time (i.e., month) fixed effects. The mean spreads across the sample are also listed in the table. The sample period underlying this regression ranges from 2012 to 2020.

	91-30	182-91	273-182	365-273	365-30
$S_{m,t-1}$	-0.0109^{***}	-0.0064^{***}	0.0005	-0.0007	-0.0175^{***}
	[-4.13]	[-3.86]	[0.55]	[-0.76]	[-4.90]
$\text{Flow}_{m,t-1}$	-0.0038	-0.0011	-0.0001	0.0005	-0.0045
	[-1.64]	[-0.70]	[-0.14]	[0.52]	[-1.39]
$\text{Ret}_{m,t-1}$	0.0052	0.0028	-0.0002	0.0001	0.0079
	[0.95]	[1.03]	[-0.11]	[0.09]	[1.19]
$\text{Bid-Ask}_{m,t-1}$	-0.0164^{***}	-0.0176^{**}	-0.0019	0.0024	-0.0335^{***}
	[-3.05]	[-2.39]	[-1.53]	[1.39]	[-2.81]
Date FE	Yes	Yes	Yes	Yes	Yes
Mean Spread	0.1181	0.0201	-0.0096	-0.0040	0.1246
Observations	1,362	1,362	1,362	1,362	1,362
R^2	0.3159	0.2472	0.2133	0.2312	0.3917

A Online appendix

A.1 Additional tables and figures

Table 6: The table presents summary statistics for the relative importance of systematic disagreement ($S_{m,t}$ from equation (41)) for each ETF m in our sample. In particular, the table reports the mean, median, and standard deviation of the relative disagreement measure for each ETF, as well as the 25^{th} and 75^{th} percentile of this measure. The sample period ranges from January 2012 to December 2020.

ETF	Style	Mean	Std	$p25$	Median	$p75$
SPY	Market	0.23	0.13	0.11	0.20	0.32
DIA	Dow Jones	0.20	0.10	0.14	0.18	0.24
QQQ	Nasdaq	0.14	0.12	0.08	0.10	0.14
XLK	Technology	0.20	0.05	0.16	0.20	0.23
XLB	Materials	0.24	0.08	0.17	0.24	0.28
XLE	Energy	0.43	0.21	0.27	0.42	0.51
XLF	Financial	0.27	0.16	0.15	0.22	0.33
XLI	Industrials	0.27	0.16	0.15	0.22	0.33
XLP	Cons. Staples	0.23	0.07	0.18	0.23	0.27
XLU	Utilities	0.22	0.06	0.18	0.22	0.24
XLV	Health care	0.16	0.06	0.12	0.15	0.20
XLY	Cons. Discretionary	0.12	0.06	0.07	0.11	0.13
XOP	Oil & gas	0.28	0.21	0.14	0.19	0.34
XBI	Biotech	0.15	0.17	0.06	0.08	0.14

Table 7: The table presents the correlation between the measures of the relative importance of systematic disagreement ($S_{m,t}$ from equation (41)) for each pair of ETFs in our sample. The sample period ranges from January 2012 to December 2020.

	SPY	DIA	QQQ	XLK	XLB	XLE	XLF	XLI	XLP	XLU	XLV	XLY	XOP	XBI
SPY	1.00	0.35	-0.08	0.05	-0.11	0.21	-0.10	-0.43	0.81	0.10	-0.17	0.46	0.02	0.10
DIA	-	1.00	-0.13	-0.08	0.20	0.29	-0.04	-0.31	0.32	0.64	-0.16	0.54	-0.03	0.64
QQQ	-	-	1.00	0.33	0.03	-0.02	0.50	0.23	-0.16	0.05	0.08	-0.23	-0.11	0.05
XLK	-	-	-	1.00	0.10	-0.24	-0.04	0.09	-0.10	-0.16	0.09	-0.25	-0.07	-0.16
XLB	-	-	-	-	1.00	0.16	0.19	0.01	-0.25	0.02	-0.07	-0.05	0.18	0.02
XLE	-	-	-	-	-	1.00	0.06	-0.41	0.16	0.20	-0.23	0.42	0.23	0.20
XLF	-	-	-	-	-	-	1.00	0.15	-0.09	0.16	0.07	0.08	0.02	0.16
XLI	-	-	-	-	-	-	-	1.00	-0.27	0.00	0.40	-0.36	0.16	0.00
XLP	-	-	-	-	-	-	-	-	1.00	0.29	0.04	0.59	0.03	0.29
XLU	-	-	-	-	-	-	-	-	-	1.00	0.27	0.49	0.04	1.00
XLV	-	-	-	-	-	-	-	-	-	-	1.00	-0.17	-0.02	0.27
XLY	-	-	-	-	-	-	-	-	-	-	-	1.00	0.24	0.49
XOP	-	-	-	-	-	-	-	-	-	-	-	-	1.00	0.04
XBI	-	-	-	-	-	-	-	-	-	-	-	-	-	1.00