Implied Dividend Volatility and Expected Growth

By Niels J. Gormsen, Ralph S.J. Koijen, and Ian Martin*

A large literature is concerned with measuring economic uncertainty and quantifying its impact on real decisions, such as investment, hiring, and R&D, and ultimately economic growth (Bloom, 2009; Jurado et al., 2015). The COVID-19 pandemic underscores the importance of timely measures of uncertainty and expected growth across horizons.

Asset prices, such as dividend futures (van Binsbergen et al., 2013; Gormsen and Koijen, 2020) and index options (Gao and Martin, 2020), provide particularly useful measures as they are forward looking and available at high frequencies. Dividend futures are claims on the dividends of the aggregate stock market in a particular year. As dividend futures are differentiated by maturity, just like nominal and real bonds, we can use these prices to obtain growth expectations by maturity.

We extend this literature by using new data on the prices of options on index-level dividends, from which we can compute *implied dividend volatility*. These implied volatilities differ from the VIX which measures uncertainty about stock prices, not only uncertainty about dividends.

We construct a term structure of implied dividend volatilities that characterizes how uncertainty varies across horizons. We study how this term structure developed over the COVID-19 crisis, documenting a substantial increase in the volatility of near-future dividends that lingers even as the volatility of the overall market portfolio has started to fall.

In addition to introducing this market, we also provide new theoretical results that show how these data can be used to derive lower bounds on expected dividend returns and on expected growth rates, by maturity, by exploiting the insights of Martin (2017). This provides an alternative to methods used in the literature using vector autoregressions or survey expectations, and sharpens alternative bounds in the literature.

I. Pricing and Riskiness of Dividends

We denote dividends on the aggregate stock market by D_t and S_t denotes the value of the aggregate stock market. The present-value identity implies

$$S_t = \sum_{\tau=1}^{\infty} E_t \left[D_{t+\tau} M_{t:t+\tau} \right], \tag{1}$$

where $M_{t:t+\tau} = \prod_{s=1}^{\tau} M_{t+s}$ and M_t denotes the stochastic discount factor. We define the τ-period dividend strip as $P_t(\tau) = E_t \left[D_{t+\tau} M_{t:t+\tau} \right]$ and the dividend futures price as $F_t(\tau) = P_t(\tau)/E_t \left[M_{t:t+\tau} \right]$.

A. Data

We use data on dividend futures for the S&P 500 index in the US (SPX) and for the Euro Stoxx 50 index in Europe (SX5E). We source these data from Bloomberg. We also use data on two ETFs; one ETF tracks long-term Treasuries (with ticker TLT) and the other ETF tracks the investment-grade corporate bond market in the US (with ticker LQD). These data are from the Center for Research in Security Prices.

We use data on Euro Stoxx 50 dividend options trading on the Eurex Exchange. These are European options on index dividends. The ten nearest successive annual contracts of the December cy-

* Gormsen: University of Chicago, Booth School of Business, 5807 South Woodlawn Avenue, 60637 Chicago, Illinois, USA, niels.gormsen@chicagobooth.edu. Koijen: University of Chicago, Booth School of Business, 5807 South Woodlawn Avenue, 60637 Chicago, Illinois, USA, ralph.koijen@chicagobooth.edu. Martin: London School of Economics, Houghton Street, London WC2A 2AE, United Kingdom, i.w.martin@lse.ac.uk. We thank Lasse Pedersen for comments and suggestions.

cle are available for trading at any point in time. The Eurex Exchange records the daily settlement prices on the options and computes the ATM implied volatilities using a Black (76) model. We source these volatilities from Bloomberg. We note that liquidity in the dividend options market is a concern. We view the current paper mostly as a proof of concept of what can be learned from these markets that may be more widely traded in the future.

For dividend options, the maturity coincides with the year in which the dividends are paid. This implies that we simultaneously vary the timing of the dividend and the maturity of the option. We use the December 2021, December 2022, and December 2023 contracts in our analysis.[1] For the Euro Stoxx 50, we choose the 12-month and 24-month implied volatilities (VSTOXX), which we get from Bloomberg. We linearly interpolate both series to target the December 2021 maturity for the market as well. All volatilities are annualized.

We sample our data weekly and use a sample from January 2020 until October 2020.

B. *Empirical results*

We study how prices and implied volatilities of both indexes and dividends changed during the COVID-19 crisis.

Figure B in the Online Appendix shows that aggregate stock markets in Europe and the US fell by 20-30%, while short-term dividends fell even more for both indexes. During the same period, Treasuries rallied and investment-grade corporate bonds fell by about 10%.

Financial markets recovered since then, with the US stock market and the investment-grade corporate bond market recovering fully and the European stock market recovering about half of its losses. However, short-term dividends have not experienced the same recovery. Prices are still down by almost 20% in the US and more than 30% in Europe, suggesting that the market prices substantial losses in the near term.

The left panel of Figure I.B shows the implied dividend volatility of the December 2021 option and the implied volatility of the aggregate stock market. In the right panel, we plot the term structure of implied dividend volatilities for the 2020, 2021, and 2022 contracts, alongside the implied volatility of the market, in January, March, and October of 2020.

Implied volatilities before the COVID-19 crisis increase with maturity, and are particularly low for the 2021 contract. The level of volatility is comparable to the historical annual dividend volatility.

During the crisis, the implied dividend volatility increases sharply and, in case of the 2021 contract, rises above the implied volatility of the market. This increase shows that short-term dividend growth is strongly heteroskedastic. The volatility of the short-term dividends remains high at the end of our sample, with the volatility of the short-term claims approximately at the same level as the market.

As such, the relative increase in volatility during the crisis is much stronger for the short-term dividends than for the market and the increase is more persistent. A key takeaway from this section is that financial markets price the pandemic via lower dividend prices and high uncertainty about short-term cash flows. As such, while the market indexes have largely recovered, the pandemic is still reflected in the pricing of near-future cash flows.

II. Expected Returns and Growth

The price of a dividend claim reflects a combination of the expected return on the claim and the expected dividend. We will derive a lower bound on the expected return, and hence on the expected dividend.

A. *Methodology*

We define $R^{\tau}_{t+\tau}$ as the spot-return on the τ-period dividend claim, $R^{\tau}_{t+\tau} = \frac{D_{t+\tau}}{P_t(\tau)}$. Our starting point is the following identity,

[1] While the December 2020 contract is also available, its implied volatility mechanically dwindles during our sample as more dividends are announced.

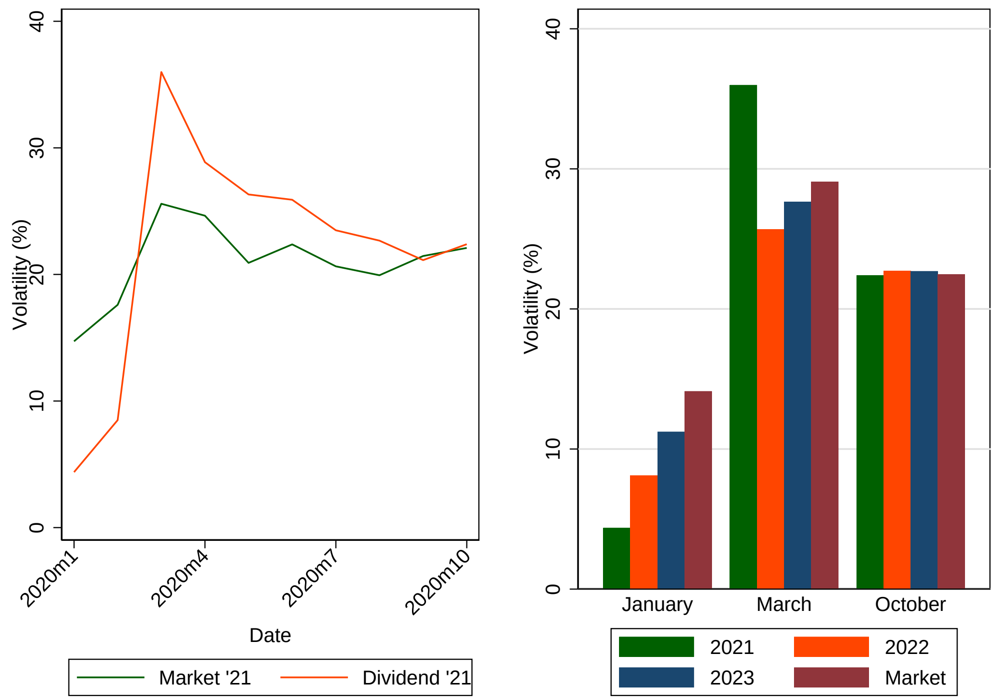

FIGURE 1. VOLATILITY DYNAMICS DURING THE COVID-19 CRISIS. THE YEARS IN THE RIGHT PANEL CORRESPOND TO THE MATURITIES OF THE IMPLIED DIVIDEND VOLATILITIES.

which holds for any returns $R_{t+\tau}$ and $R^{\tau}_{t+\tau}$:

$$E_t[R^{\tau}_{t+\tau}] - R^f_t = \frac{\text{cov}^*_t(R^{\tau}_{t+\tau}, R_{t+\tau})}{R^f_t} - \text{cov}_t(M_{t+\tau}R_{t+\tau}, R^{\tau}_{t+\tau}).$$

We use asterisks to denote risk-neutral moments and write R^f_t for the gross risk-free rate between t and $t+\tau$. This relationship—a variant of the identity in Martin (2017)—was exploited in the context of currency returns by Kremens and Martin (2019).

We explore two sets of assumptions to use this identity to derive a lower bound for expected returns and expected dividend growth rates. We will refer to these approaches as Method 1 and Method 2.

Method 1: We pick $R_{t+\tau} = R^{\tau}_{t+\tau}$, which means that the first term on the RHS is the risk-neutral variance on the dividend and thus observable. We assume that the second covariance term is non-positive. If the return on the dividend claim is proportional to the asset return (as is the case in models that generate constant price-dividend ratios) then this condition reduces to the NCC of Martin (2017). This applies, for example, to models such as Barro (2006), in which the NCC holds. Accordingly, the risk-neutral variance of $R_{t+\tau}(\tau)$ constitutes a lower bound on expected returns, giving rise to the following bound on expected dividends:

$$E_t[D_{t+\tau}] \geq P_t(\tau)\left(\frac{\text{var}^*_t(R^{\tau}_{t+\tau})}{R^f_t} + R^f_t\right)$$

In Appendix A, we show that the covariance is negative in bad states of the world in the Campbell and Cochrane (1999) external habit model, but is positive in the Bansal and Yaron (2004) long-run risks model for short maturities. Also, the covariance term is potentially positive in models in which returns on dividends are not strongly negatively correlated with the SDF. We therefore also consider an alternative assumption.

Method 2: We (i) pick $R_t = R_t^M$, which is the return on the aggregate stock market, and (ii) assume that returns, dividends, and the SDF are jointly log-normally distributed.[2] As we show in Appendix A, this implies

$$E_t[R_{t+\tau}^{\tau}] - R_t^f = \frac{\rho_t \sigma_t^*(R_{t+\tau}^{\tau}) \sigma_t^*(R_{t+\tau}^M)}{R_t^f} - \text{cov}_t(M_{t+\tau} R_{t+\tau}^M, R_{t+\tau}^{\tau}),$$

where $\rho_t = \text{corr}_t(R_{t+\tau}^{\tau}, R_{t+\tau}^M)$ and σ_t^* denotes risk-neutral volatility. We further assume that the second covariance term is non-positive. (If one adopts the perspective of an investor with log utility who chooses to invest fully in the market, then $M_{t+\tau} = 1/R_{t+\tau}^M$ so that the covariance term is exactly zero. In this case, the bound holds with equality.) The lower bound above again gives rise to a lower bound on expected dividends:

$$E_t[D_{t+\tau}] \geq P_t(\tau) \left(\frac{\rho_t \sigma_t^*(R_{t+\tau}^{\tau}) \sigma_t^*(R_{t+\tau}^M)}{R_t^f} + R_t^f \right)$$

Method 2 requires an estimate of $\text{corr}_t(R_{t+\tau}^{\tau}, R_{t+\tau}^M)$. As estimating a correlation model is beyond the scope of this paper, we will present results for a range of values. We remark that the assumption of log-normality is likely violated, as discussed by Martin (2017), but the hope is that this violation has limited impact on the final results. We leave it for future research to derive bounds under more general distributional assumptions.

B. *Empirical Results*

We approximate $R_t^f \simeq 1$ during our sample. Method 1 implies that the lower bound on the annualized expected excess returns on the 2021 dividend varies from 0.2% before the crisis to 12.4% during the downturn, and 5.0% at the end of our sample.

This points to high expected excess returns on short-term claims during the crisis, consistent with van Binsbergen et al. (2012) and van Binsbergen et al. (2013). However, it is inconsistent with the models that motivate the covariance constraint in the first place. Hence, we are less comfortable using this bound, and these results are best viewed as a rejection of the models in which the conditional covariance condition is satisfied as the implications for dividend volatility under the null of the model lead to too much volatility in discount rates on short-term claims.

Method 2 instead relies on a covariance restriction that mimics Martin (2017). As estimating the conditional correlation is challenging, we consider two values that we consider plausible during times of stress, $\rho_t = 50\%$ or 75%.

For these values, we plot the lower bound for expected excess returns in annualized terms in the left panel of Figure II.B. The expected excess return peaks at 4.5%, when $\rho_t = 50\%$, or 6.7%, when $\rho_t = 75\%$. As before, this suggests substantial expected excess returns on dividend claims.[3]

In the right panel of Figure II.B, we use these estimates to compute a lower bound on the expected dividend in December 2021. To simplify the interpretation, we scale this expectation by the December 2019 dividend. This bound sharpens the lower bound in Gormsen and Koijen (2020), which corresponds to $\rho_t = 0$, and provides an alternative to VAR methods and survey expectations.

We draw two conclusions from this figure. First, even though we tighten the lower bound—as the lower bound for expected excess returns is well above zero—most of the variation in dividend futures is due to growth expectations. This underscores the usefulness of dividend futures for

[2]We need the latter assumption because derivatives whose prices would reveal the risk-neutral covariance between the market return and dividend growth are not widely traded. By contrast, Kremens and Martin (2019) were able to exploit the fact that index quanto contracts, which reveal the corresponding risk-neutral covariance between the market return and currency appreciation, are traded.

[3]As the volatilities of all dividend claims and the market move significantly during the crisis, so do the lower bounds on the expected excess returns. We refer to Gormsen (2020) for estimates of variation in expected excess returns across maturities.

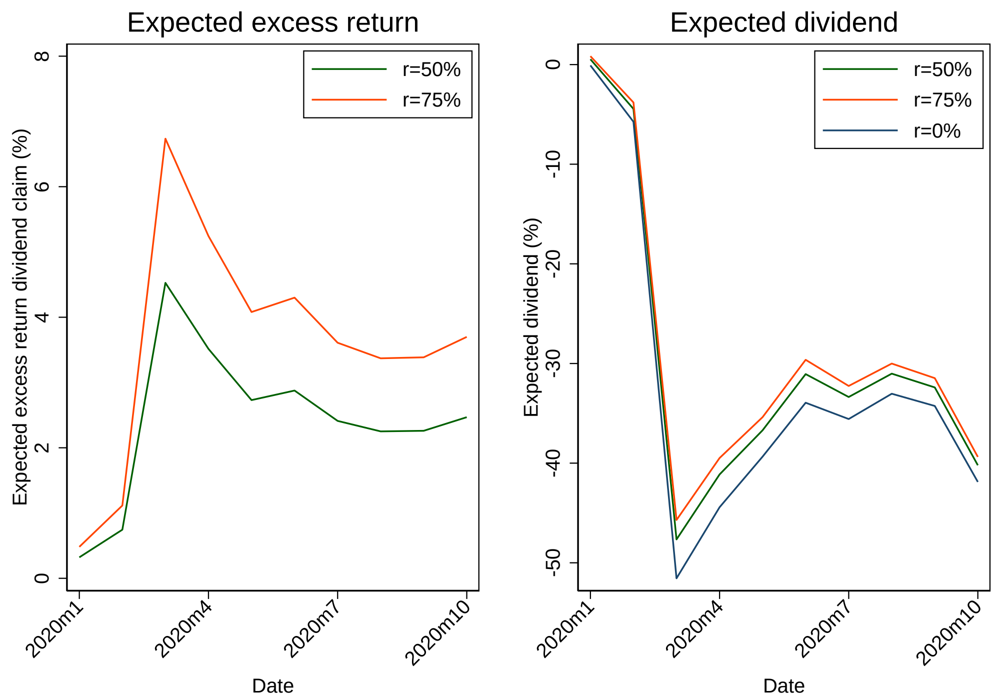

Figure 2. A lower bound on expected excess returns on the short-term dividend claim (left panel)and on the expected dividend in 2021, scaled by the 2019 dividend (right panel). The different lines correspond to conditional correlations between market returns and short-term dividend returns.

forecasting economic growth. Second, while the market recovers, in the second part of 2020, the near-term growth expectations improve only slightly and, in fact, deteriorate towards the year's end. Paired with the high level of implied dividend volatilities, the short-term economic outlook is uncertain and not expected to recover in the near term in Europe.

REFERENCES

Ravi Bansal and Amir Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance*, 59 (4):1481–1509, 2004.

Ravi Bansal, Dana Kiku, and Amir Yaron. An empirical evaluation of the long-run risks model for asset prices. *Critical Finance Review*, 1:183–221, 2012.

Robert J Barro. Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, 121(3): 823–866, 2006.

Nicholas Bloom. The impact of uncertainty shocks. *Econometrica*, 77(3):623–685, 2009.

Tim Bollerslev, George Tauchen, and Hao Zhou. Expected stock returns and variance risk premia. *Review of Financial Studies*, 22:4463–4493, 2009.

John Y. Campbell and John H. Cochrane. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251, 1999.

Ian Dew-Becker, Stefano Giglio, Anh Le, and Marius Rodriguez. The price of variance risk. *Journal of Financial Economics*, 123(2):225 – 250, 2017. ISSN 0304-405X. doi: https://doi.org/10.1016/j.jfineco.2016.04.003.

Itamar Drechsler and Amir Yaron. What's vol got to do with it. *Review of Financial Studies*, 24:1–45, 2011.

Can Gao and Ian Martin. Volatility, Valuation Ratios, and Bubbles: An Empirical Measure of Market Sentiment. *working paper*, 2020.

Niels Joachim Gormsen. Time variation of the equity term structure. *Journal of Finance, forthcoming*, 2020.

Niels Joachim Gormsen and Ralph S J Koijen. Coronavirus: Impact on Stock Prices and Growth Expectations. *The Review of Asset Pricing Studies*, 10(4):574–597, 2020.

Kyle Jurado, Sydney C. Ludvigson, and Serena Ng. Measuring uncertainty. *American Economic Review*, 105(3): 1177–1216, March 2015.

Lukas Kremens and Ian Martin. The Quanto Theory of Exchange Rates. *American Economic Review*, 109(3):810–843, 2019.

Ian Martin. What is the Expected Return on the Market? *The Quarterly Journal of Economics*, 132(1):367–433, 2017.

Jules van Binsbergen, Michael Brandt, and Ralph Koijen. On the timing and pricing of dividends. *American Economic Review*, 102(4):1596–1618, 2012.

Jules H. van Binsbergen, Wouter Hueskes, Ralph S.J. Koijen, and Evert B. Vrugt. Equity yields. *Journal of Financial Economics*, 110(3):503 – 519, 2013.

Jessica A. Wachter. Can time-varying risk of rare disasters explain aggregate stock market volatility? *Journal of Finance*, 68:987–1035, 2013.

Online appendix

A1. Method 1

We evaluate whether the covariance condition holds such that the bound is indeed a lower bound in several asset pricing models.

Lucas model

Consider a Lucas model with power utility and similar technology processes as in Campbell and Cochrane (1999)

$$\begin{aligned}\Delta c_{t+1} &= g + v_{t+1},\\ \Delta d_{t+1} &= g + w_{t+1},\end{aligned}$$

where $v_{t+1} \sim N(0, \sigma^2)$, $w_{t+1} \sim N(0, \sigma_w^2)$, and $corr(v_{t+1}, w_{t+1}) = \rho$. The dividend price is

$$\begin{aligned}\mathbb{E}_t\left(M_{t+1}D_{t+1}\right) &= \delta D_t\, \mathbb{E}_t\left(\exp\left\{(1-\gamma)g - \gamma v_{t+1} + w_{t+1}\right\}\right)\\ &= D_t R_f^{-1} \exp\left\{g + \frac{1}{2}\sigma_w^2 - \gamma\sigma\sigma_w\rho\right\},\end{aligned}$$

where $R_f = \delta^{-1}\exp\{\gamma g - \frac{1}{2}\gamma^2\sigma^2\}$. The dividend return, $R_{d,t+1} = D_{t+1}/\,\mathbb{E}_t\left(M_{t+1}D_{t+1}\right)$, is

$$R_{d,t+1} = R_f \exp\left\{\gamma\sigma\sigma_w\rho - \frac{1}{2}\sigma_w^2 + w_{t+1}\right\},$$

so $\mathbb{E}\,R_{d,t+1} = R_f \exp\{\gamma\sigma\sigma_w\rho\}$ and $\operatorname{var}\log R_{d,t+1} = \sigma_w^2$. As M_{t+1} and $R_{d,t+1}$ are conditionally lognormal, the NCC holds for the dividend strip return if and only if

$$\frac{\log \mathbb{E}_t\, R_{d,t+1} - \log R_f}{\sigma_t(\log R_{d,t+1})} \geq \sigma_t(\log R_{d,t+1}).$$

The quantity on the left-hand side of the inequality is, essentially, the Sharpe ratio of the dividend strip. Hence the NCC condition holds if and only if

$$\gamma\rho\sigma \geq \sigma_w.$$

However, we can also compute the covariance term and avoid bounds altogether

$$\operatorname{cov}\left(M_{t+1}R_{d,t+1}, R_{d,t+1}\right) = \mathbb{E}\left(M_{t+1}R_{d,t+1}^2\right) - \mathbb{E}\left(R_{d,t+1}\right).$$

The first term is given by

$$\begin{aligned}\mathbb{E}\left(M_{t+1}R_{d,t+1}^2\right) &= \mathbb{E}\left(R_f^2 \exp\left\{2\gamma\sigma\sigma_w\rho - \sigma_w^2 + 2w_{t+1} - \gamma g - \gamma v_{t+1}\right\}\right)\\ &= R_f^2 \exp\left\{2\gamma\sigma\sigma_w\rho - \sigma_w^2 + 2\sigma_w^2 - \gamma g + \frac{1}{2}\gamma^2\sigma^2 - 2\sigma\sigma_w\rho\gamma\right\}\\ &= R_f \exp\left\{\sigma_w^2\right\}\end{aligned}$$

Putting it all together implies that

$$\begin{aligned}\operatorname{var}^*\left(R_{d,t+1}\right) &= R_f\left[\mathbb{E}\,R_{d,t+1} - R_f + \operatorname{cov}\left(M_{t+1}R_{d,t+1}, R_{d,t+1}\right)\right]\\ &= R_f^2\left(\exp\{\sigma_w^2\} - 1\right).\end{aligned}$$

We observe the left-hand side and notice that it is highly volatile, which rejects the model. More broadly, in any model of the form

$$\begin{aligned} M_{t+1} &= R_f^{-1} \exp\{-\frac{1}{2}\lambda_t^2\sigma^2 - \lambda_t v_{t+1}\}, \\ R_{d,t+1} &= R_f \exp\{\mu_t - \frac{1}{2}\sigma_w + w_{t+1}\}, \end{aligned}$$

for any risk price λ_t and $\mu_t = \rho\lambda_t\sigma\sigma_w$, it holds

$$\text{var}^*\left(R_{d,t+1}\right) = R_f^2\left(\exp\{\sigma_w^2\} - 1\right).$$

These conditions are satisfied in the Campbell and Cochrane (1999) model to which we turn next.

Campbell and Cochrane (1999)

In this subsection, we use the notation from Campbell and Cochrane (1999) without further comment. The price of a claim to the first dividend is

$$\begin{aligned} \mathbb{E}_t\left(M_{t+1}D_{t+1}\right) &= D_t\,\mathbb{E}_t\left(M_{t+1}\frac{D_{t+1}}{D_t}\right) \\ &= \delta D_t\,\mathbb{E}_t\left(\exp\left\{-\gamma g - \gamma\left\{(\phi-1)(s_t-\overline{s}) + [1+\lambda(s_t)]v_{t+1}\right\} + g + w_{t+1}\right\}\right) \\ &= \delta D_t \exp\left\{g(1-\gamma) + \gamma(1-\phi)(s_t-\overline{s}) + \frac{\gamma^2\sigma^2}{2\overline{S}^2}(1-2(s_t-\overline{s})) + \frac{1}{2}\sigma_w^2 - \frac{\gamma\rho\sigma\sigma_w}{\overline{S}}\sqrt{1-2(s_t-\overline{s})}\right\}. \end{aligned}$$

So the return on this dividend strip is

$$R_{d,t+1} = \frac{D_{t+1}}{\mathbb{E}_t\left(M_{t+1}D_{t+1}\right)} = \frac{D_{t+1}}{D_t}\delta^{-1}\exp\left\{g(\gamma-1) - \gamma(1-\phi)(s_t-\overline{s}) - \frac{\gamma^2\sigma^2}{2\overline{S}^2}(1-2(s_t-\overline{s})) - \frac{1}{2}\sigma_w^2 + \frac{\gamma\rho\sigma\sigma_w}{\overline{S}}\sqrt{1-2(s_t-\overline{s})}\right\}.$$

Hence the expected return is

$$\begin{aligned} \mathbb{E}_t\,R_{d,t+1} &= \delta^{-1}\exp\left\{g\gamma - \gamma(1-\phi)(s_t-\overline{s}) - \frac{\gamma^2\sigma^2}{2\overline{S}^2}(1-2(s_t-\overline{s})) + \frac{\gamma\rho\sigma\sigma_w}{\overline{S}}\sqrt{1-2(s_t-\overline{s})}\right\} \\ &= R_f \exp\left\{\frac{\gamma\rho\sigma\sigma_w}{\overline{S}}\sqrt{1-2(s_t-\overline{s})}\right\}, \end{aligned}$$

and

$$\text{var}_t \log R_{d,t+1} = \sigma_w^2.$$

From the above equations, the NCC holds for the dividend strip return if and only if

$$\frac{\gamma\rho\sigma}{\overline{S}}\sqrt{1-2(s_t-\overline{s})} \geq \sigma_w.$$

This condition mimics the Lucas model above when $s_t - \overline{s}$ and $\gamma^{Lucas} = \frac{\gamma}{\overline{S}}$.

Figure A.A1 shows that the NCC holds in sufficiently bad states of the world, but not in good states of the world and not at the steady state level of habit, $\overline{S}$.

Bansal and Yaron (2004)

In this subsection, we use the notation from Bansal and Yaron (2004) without further comment. We focus on the Case II calibration that features stochastic volatility.

Single-period calculations. A claim to the first dividend earns zero risk premium because dividend growth is conditionally uncorrelated with the log stochastic discount fac-

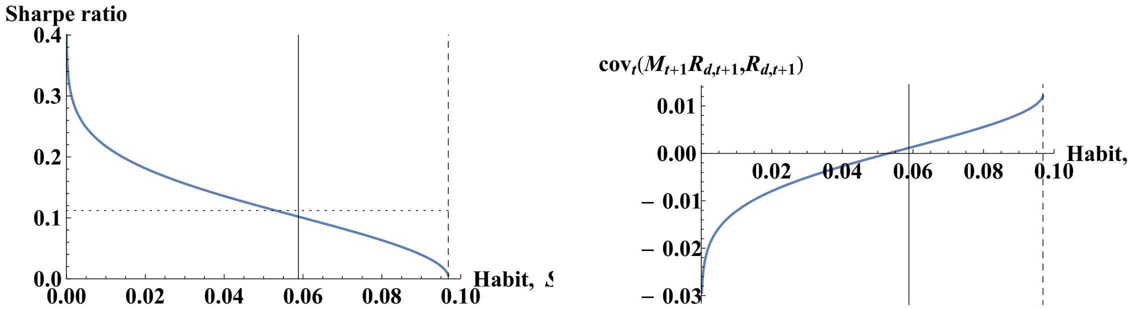

Figure A1. The Sharpe ratio, and the covariance term, for the one-period dividend strip in Campbell and Cochrane (1999). The dotted line in the left panel indicates the volatility of the dividend strip return, σ_w. The solid vertical line indicates the steady state level of habit. The dashed vertical line indicates the maximum attainable level of habit.

Note: Figure notes without optional leading.

tor, $\text{cov}_t(g_{d,t+1}, m_{t+1}) = 0$. (We are following Bansal and Yaron by treating the model as conditionally lognormal, relying on loglinearizations.) Hence $\mathbb{E}_t R_{d,t+1} = R_{f,t+1}$. Again exploiting lognormality, the risk-neutral variance takes the form above,

$$\text{var}_t^* R_{d,t+1} = R_{f,t+1}^2 \left(e^{\text{var}_t \log g_{d,t+1}} - 1\right) = R_{f,t+1}^2 \left(e^{\varphi_d^2 \sigma_t^2} - 1\right),$$

and

$$\text{cov}_t \left(M_{t+1} R_{d,t+1}, R_{d,t+1}\right) = \frac{1}{R_{f,t+1}} \text{var}_t^* R_{d,t+1} - \left(\mathbb{E}_t R_{d,t+1} - R_{f,t+1}\right) = R_{f,t+1} \left(e^{\varphi_d^2 \sigma_t^2} - 1\right).$$

Hence the conditional covariance is positive (but very small) in Bansal and Yaron (2004).

Bansal et al. (2012) make dividends and consumption growth correlated in the short run. As a result, $\text{cov}_t(g_{d,t+1}, m_{t+1}) = -\gamma\pi\sigma_t^2$, and hence $\mathbb{E}_t R_{d,t+1} = R_{f,t+1} e^{\gamma\pi\sigma_t^2}$, where π is a new parameter introduced in equation (3) of Bansal et al. (2012). Risk-neutral variance changes slightly:

$$\text{var}_t^* R_{d,t+1} = R_{f,t+1}^2 \left(e^{\left(\pi^2+\varphi^2\right)\sigma_t^2} - 1\right)$$

and

$$\text{cov}_t \left(M_{t+1} R_{d,t+1}, R_{d,t+1}\right) = R_{f,t+1} \left(e^{\left(\pi^2+\varphi^2\right)\sigma_t^2} - e^{\gamma\pi\sigma_t^2}\right).$$

This is positive in their calibration, in which $\pi^2 + \varphi^2 = 42.3 > \gamma\pi = 26$.

Multi-period calculations. Bansal and Yaron (2004) calibrate the model to a monthly frequency, while we use 1- to 3-year dividend claims. We therefore compute the risk-neutral variance and the covariance for longer horizons for completeness. The model implies that

$$P_t(\tau) = D_t \exp(A_\tau + B_\tau x_t + C_\tau \sigma_t^2),$$

where the coefficients follow from

$$\begin{aligned} P_t(\tau) &= D_t \, \mathbb{E}_t \left(M_{t+1} \exp\left\{g_{d,t+1} + A_{\tau-1} + B_{\tau-1} x_{t+1} + C_{\tau-1} \sigma_{t+1}^2\right\}\right) \\ &= D_t \exp\left\{\mu_d + \phi x_t + A_{\tau-1} + B_{\tau-1} \rho x_t + C_{\tau-1} \sigma^2 (1-\nu_1) + C_{\tau-1} \nu_1 \sigma_t^2\right\} \times \\ & \quad R_{ft}^{-1} \exp\left\{\frac{1}{2}\left(\varphi_d^2 + B_{\tau-1}^2 \varphi_e^2\right) \sigma_t^2 + \frac{1}{2} C_{\tau-1}^2 \sigma_w^2 - \lambda_{m,e} B_{\tau-1} \varphi_e \sigma_t^2 - \lambda_{m,w} C_{\tau-1} \sigma_w^2\right\} \end{aligned}$$

where $R_{ft} = \exp\{s_0 + s_1 x_t + s_2 \sigma_t^2\}$ and thus

$$\begin{aligned} A_\tau &= -s_0 + \mu_d + A_{\tau-1} + C_{\tau-1}\sigma^2(1-\nu_1) + \frac{1}{2}C_{\tau-1}^2\sigma_w^2 - \lambda_{m,w}C_{\tau-1}\sigma_w^2, \\ B_\tau &= -s_1 + \phi + B_{\tau-1}\rho, \\ C_\tau &= -s_2 + C_{\tau-1}\nu_1 + \frac{1}{2}\left(\varphi_d^2 + B_{\tau-1}^2\varphi_e^2\right) - \lambda_{m,e}B_{\tau-1}\varphi_e. \end{aligned}$$

Returns are given by

$$\begin{aligned} R_{d,t+1}(\tau) &= \exp\left\{g_{d,t+1} + A_{\tau-1} - A_\tau + B_{\tau-1}x_{t+1} - B_\tau x_t + C_{\tau-1}\sigma_{t+1}^2 - C_\tau\sigma_t^2\right\} \\ &= f_t(\tau, x_t, \sigma_t^2)\exp\left\{\varphi_d\sigma_t u_{t+1} + B_{\tau-1}\varphi_e\sigma_t e_{t+1} + C_{\tau-1}\sigma_w w_{t+1}\right\}, \end{aligned}$$

where

$$\begin{aligned} f_t(\tau, x_t, \sigma_t^2) &= \exp\left\{\mu_d + \phi x_t + A_{\tau-1} - A_\tau + (B_{\tau-1}\rho - B_\tau)x_t + C_{\tau-1}\sigma^2(1-\nu_1) + (C_{\tau-1}\nu_1 - C_\tau)\sigma_t^2\right\} \\ &= R_{ft}\exp\left\{-\frac{1}{2}C_{\tau-1}^2\sigma_w^2 - \frac{1}{2}\left(\varphi_d^2 + B_{\tau-1}^2\varphi_e^2\right)\sigma_t^2 + \lambda_{m,w}C_{\tau-1}\sigma_w^2 + \lambda_{m,e}B_{\tau-1}\varphi_e\sigma_t^2\right\}. \end{aligned}$$

Exploiting lognormality,

$$\mathrm{var}_t^* R_{d,t+1} = R_{ft}^2\left(\exp\left\{\mathrm{var}_t \log R_{d,t+1}(\tau)\right\} - 1\right) = R_{ft}^2\left(\exp\left\{\left(\varphi_d^2 + B_{\tau-1}^2\varphi_e^2\right)\sigma_t^2 + C_{\tau-1}^2\sigma_w^2\right\} - 1\right).$$

For the risk premium, we have

$$\mathbb{E}_t\left(R_{d,t+1}(\tau)\right) - R_{ft} = R_{ft}\left(\exp\{\lambda_{m,w}C_{\tau-1}\sigma_w^2 + \lambda_{m,e}B_{\tau-1}\varphi_e\sigma_t^2\} - 1\right).$$

This implies for the covariance

$$\begin{aligned} \mathrm{cov}_t\left(M_{t+1}R_{d,t+1}(\tau), R_{d,t+1}(\tau)\right) = {} & R_{ft}(\exp\{\left(\varphi_d^2 + B_{\tau-1}^2\varphi_e^2\right)\sigma_t^2 + C_{\tau-1}^2\sigma_w^2\} \\ & - \exp\{\lambda_{m,w}C_{\tau-1}\sigma_w^2 + \lambda_{m,e}B_{\tau-1}\varphi_e\sigma_t^2\}), \end{aligned}$$

which we linearize (and approximating $R_{ft} = 1$ as it does not affect the sign and is the relevant empirical case)

$$\mathrm{cov}_t\left(M_{t+1}R_{d,t+1}(\tau), R_{d,t+1}(\tau)\right) \simeq \varphi_d^2\sigma_t^2 + B_{\tau-1}\varphi_e\left(B_{\tau-1}\varphi_e - \lambda_{m,e}\right)\sigma_t^2 + C_{\tau-1}\left(C_{\tau-1} - \lambda_{m,w}\right)\sigma_w^2.$$

Note that $B_\tau, B'_\tau > 0$, $C_\tau, C'_\tau < 0$, for $\tau > 0$, and $\lambda_{m,e} > 0$ and $\lambda_{m,w} < 0$. At longer horizons, the NCC will be satisfied, but the coefficients (B_τ, C_τ) change only slowly with maturity due to the persistence of the processes. We therefore conclude that the NCC condition of method 1 is likely not satisfied in Bansal and Yaron (2004) when calibrated to our sample period.

A2. Method 2

Additional calculations under lognormality

If $M_{t+1} = e^{-r_{f,t} - \frac{1}{2}\sigma_{1,t}^2 - \sigma_{1,t}Z_{1,t+1}}$ and $R_{d,t+1} = e^{\mu_{d,t} - \frac{1}{2}\sigma_{2,t}^2 + \sigma_{2,t}Z_{2,t+1}}$ and $R_{t+1} = e^{\mu_t - \frac{1}{2}\sigma_{3,t}^2 + \sigma_{3,t}Z_{3,t+1}}$ then we must have $\mu_{d,t} - r_{f,t} = \rho_{12,t}\sigma_{1,t}\sigma_{2,t}$ and $\mu_t - r_{f,t} = \rho_{13,t}\sigma_{1,t}\sigma_{3,t}$ so that $\mathbb{E}\, MR = 1$ holds, where we are writing $\rho_{ij,t}$ for $\mathrm{corr}_t(Z_{i,t+1}, Z_{j,t+1}$. Straightforward calculations show that $\mathrm{cov}_t(R_{d,t+1}, R_{t+1}) = e^{\mu_{d,t}+\mu_t}(e^{\rho_{23,t}\sigma_{2,t}\sigma_{3,t}} - 1)$ and $\mathrm{cov}_t^*(R_{d,t+1}, R_{t+1}) = e^{2r_{f,t}}(e^{\rho_{23,t}\sigma_{2,t}\sigma_{3,t}} - 1)$; similarly, $\mathrm{var}_t R_{d,t+1} = e^{2\mu_{d,t}}(e^{\sigma_{2,t}^2} - 1)$ and $\mathrm{var}_t R_{t+1} = e^{2\mu_t}(e^{\sigma_{3,t}^2} - 1)$, while risk-neutral variances are $\mathrm{var}_t^* R_{d,t+1} = e^{2r_{f,t}}(e^{\sigma_{2,t}^2} - 1)$ and $\mathrm{var}_t^* R_{t+1} = e^{2r_{f,t}}(e^{\sigma_{3,t}^2} - 1)$.

It follows that the true and risk-neutral correlations are equal:

$$\operatorname{corr}_t\left(R_{d,t+1}, R_{t+1}\right) = \operatorname{corr}_t^*\left(R_{d,t+1}, R_{t+1}\right) = \frac{e^{\rho_{23,t}\sigma_{2,t}\sigma_{3,t}} - 1}{\sqrt{\left(e^{\sigma_{2,t}^2} - 1\right)\left(e^{\sigma_{3,t}^2} - 1\right)}} \approx \rho_{23,t}\,.$$

The covariance condition

Although we have focussed on dividend volatility, consumption-based models also have difficulty matching the time series behavior of price volatility. Martin (2017, Table IV) reports time series of various statistics (mean, median, standard deviation, min, max, skewness, kurtosis, and autocorrelation) of sample paths of the VIX and SVIX indices generated in the model economies of Campbell and Cochrane (1999), Bansal and Yaron (2004), Bansal et al. (2012), Bollerslev et al. (2009), Drechsler and Yaron (2011), and Wachter (2013). None of these is able to generate sample paths that resemble those observed empirically. All the models apart from Drechsler and Yaron (2011) generate volatility series that are more persistent than the data. The empirically observed mean gap between VIX and SVIX—a measure of the average importance of extreme left-tail events—is outside the support of the million sample paths in *every* model: Wachter (2013) overstates the importance of such events, and all other models understate it. All the models apart from Wachter (2013) fail, on 99% of sample paths, to generate the maximum levels of VIX that have been observed in reality. All models apart from Drechsler and Yaron (2011) fail, on 99% of sample paths, to match the kurtosis of VIX and SVIX.

We also refer to Dew-Becker et al. (2017) for the a set of related challenges of models to match the term structure of variance swaps and variance risk premia. Dew-Becker et al. show that the monthly risk premium on short-term variance (≤ 3 months) is negative and large whereas the risk premium on longer-term variance (> 3 months) is essentially zero. The fact that very-short run variance is priced but longer-term variance is not implies the existence of a transitory element in realized volatility that investors are highly averse towards. The results also imply that investors are not averse towards changes in long-term expected volatility – something that is hard to reconcile with long-run risk models (Drechsler and Yaron, 2011) and disaster models with Epstein-Zin preferences (Wachter, 2013) as investors in such models are averse towards increases in expected volatility, which is modeled to be persistent, meaning that claims on longer-term variance should be priced.

ADDITIONAL FIGURES

We summarize the broad patterns in the data as opposed to high-frequency event studies.[4] We therefore sample the data at three moments in time: the pre-pandemic peak of the market, the bottom of the market, and at the end of our sample. For both Europe and the US, we determine the peak of the index level before the start of the pandemic and compute the average prices and volatilities of each asset during the three week period before the peak. Similarly, we determine the bottom of the market indexes and average the prices and implied volatilities in the three weeks surrounding the bottom. We also average the prices during the last three weeks of our sample. To succinctly present the results, we present the returns averaged for the 2021, 2022, and 2023 dividend futures prices. The dividend prices are indicated by "ST" in the legend of the figure.

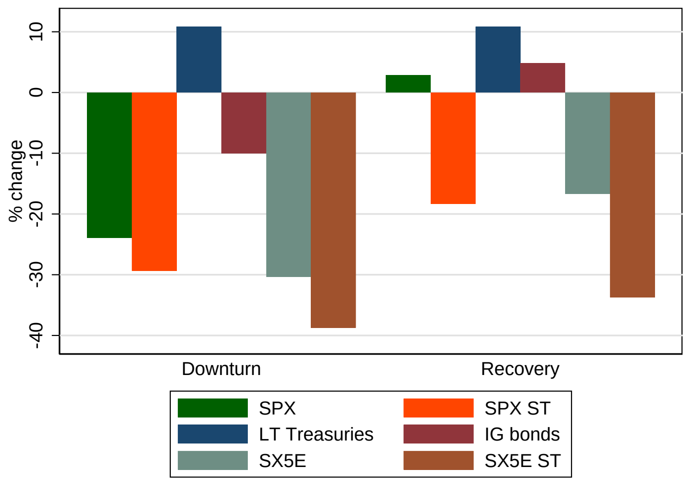

Figure B1. The dynamics of asset prices during the COVID-19 crisis.

[4]Gormsen and Koijen (2020) analyze the dynamics of the index and dividend futures around some of the key events during the crisis for the European, Japanese, and US market.