Macroeconomic Tail Risks and Asset Prices

David Schreindorfer*

September 16, 2019

Abstract

I document that dividend growth and returns on the aggregate U.S. stock market are more correlated with consumption growth in bad economic times. In a consumption-based asset pricing model with a generalized disappointment averse investor and small, IID consumption shocks, this feature results in a realistic equity premium despite low risk aversion. The model is consistent with the main facts about stock market risk premia inferred from equity index options, remains tightly parameterized, and allows for analytical solutions for asset prices. An extension with non-IID dynamics accounts for excess volatility and return predictability while preserving the model's consistency with option moments.

*W. P. Carey School of Business, Arizona State University, PO Box 873906, PO Box 873906. Email: david.schreindorfer@asu.edu. I thank Sreedhar Barath, Oliver Boguth, Mikhail Chernov, Itamar Drechsler, Lars-Alexander Kuehn, Ian Martin, Rajnish Mehra, Stefan Nagel, Stijn Van Nieuwerburgh (editor), Seth Pruitt, Bryan Routledge, and two anonymous referees for helpful comments and suggestions. The paper has also benefited from the comments and suggestions of seminar participants at Arizona State University and conference participants at the SFS Cavalcade 2018 and WFA 2018.

To paraphrase Samuelson, the option markets have predicted nine out of the past five market corrections, generating surprisingly large returns from selling crash insurance via out-of-the-money put options.

—Bates (2008)

Option contracts allow economists to infer how investors value payoffs in different states of nature. A key insight is that investors are willing to pay very large premiums for protection against market corrections. This observation is informative about potential explanations of the equity premium, because it implies that large negative stock market returns must be associated with very high marginal utility for the representative investor. While the macro finance literature has made great progress in linking risk premiums in asset markets to macroeconomic fundamentals, existing frameworks remain inconsistent with this well-known evidence from option markets. In particular, Backus et al. (2011), Martin (2017), and Dew-Becker et al. (2017) show that leading asset pricing models fail to explain option prices, option returns, or both, and Beason and Schreindorfer (2019) illustrate that the models are inconsistent with the main sources of the equity premium as a result.

This paper proposes a consumption-based explanation for risk premiums in aggregate stock markets that matches the main features of equity index options. The theory predicts that the equity premium predominantly reflects exposure to macroeconomic tail risk – drops in consumption of the magnitude typically observed in recessions.

To connect stock market tail risks to tail risks in the macroeconomy, the model relies on two key deviations from the canonical framework. First, the representative agent has generalized disappointment aversion (GDA) risk preferences, an axiomatic extension of expected utility (EU) theory that resolves the Allais Paradox and allows for asymmetric risk attitudes over gains and losses (Gul 1991; Routledge and Zin 2010). I show that GDA preferences can generate a high price for macroeconomic tail risks, while simultaneously relying on low risk aversion. This property stands in sharp contrast to EU risk preferences with low risk aversion, which imply that moderately sized

consumption shocks are close to irrelevant for risk premiums (Mehra and Prescott 1985).[1] Second, consumption and dividends follow random walks with innovations from a mixture distribution, where both growth rates are subject to independent Gaussian shocks and negatively exposed to a common exponential shock. Like the bivariate normal distribution, mixture shocks can be calibrated to match the imperfect correlation between consumption and dividends in the data. Unlike the bivariate normal distribution, however, mixture shocks result in a high likelihood of joint tail events in consumption and dividends. I show that this feature finds strong support in the data. Because bad news for equity holders (left tail events in dividends) is likely to coincide with high marginal utility for the GDA agent (left tail events in consumption), the model produces a large equity premium.

Similar to the rare disasters framework of Rietz (1988) and Barro (2006), the mechanism I propose builds on the intuitive idea that tail risks are important for understanding asset prices. The nature and magnitude of these risks, however, are quite different. The rare disaster model assumes EU risk preferences with a low risk aversion parameter. Because this assumption implies that small shocks are nearly irrelevant to investors, the equity premium depends crucially on a high probability of extreme outcomes. However, an important study by Backus et al. (2011) points out that options can be used to quantify the amount of tail risk perceived by investors, and shows that Barro's calibration to international macroeconomic data generates considerably more negative return skewness than implied by observed U.S. option prices. Beason and Schreindorfer (2019) show that this mismatch has important implications for the equity premium. In particular, they use option prices to decompose the equity premium into different parts of the return state space, and show that states associated with stock market crashes account nearly the entire equity premium in the data, but for almost none of it in the Barro (2006) model.

[1]Embedding EU risk preferences into a recursive utility framework does not alter this conclusion. For example, Bansal and Yaron (2004, table V) show that IID consumption growth shocks contribute almost nothing to the equity premium in their model.

In contrast to the rare disaster framework, my model matches the equity premium based on a small *quantity* of tail risk because (1) GDA risk preferences imply a very high *price* of tail risk and (2) mixture shocks imply that tail events in consumption and dividends are likely to coincide. I show that, as a result, the model matches a broad array of equity and equity index option moments. When fundamentals are calibrated to annual U.S. consumption and dividend data over 1930-2017 (without relying on asset pricing moments), the model replicates the average equity premium and risk-free rate, the skewness and kurtosis of monthly returns, the prices, returns, and Sharpe ratios of put options with different strikes, the difference between the VIX and SVIX indices (Martin 2017), the distribution of straddle returns (Coval and Shumway 2001), and the risk-neutral skewness and kurtosis implied by option prices. Additionally, Beason and Schreindorfer (2019) show that the model is consistent with their equity premium decomposition. Despite this quantitative success, the model remains tightly parameterized (it relies on just eight free parameters) and it yields analytical solutions for asset prices.

Perhaps the most striking property of the model is that it requires neither high risk aversion nor large consumption risks to rationalize the equity premium. Specifically, I show in Section 5.3 that the GDA agent in my calibration has the same level of risk aversion as an EU agent with a relative risk aversion coefficient of about ten. Despite the fact that consumption is only subject to moderately sized IID shocks, however, the GDA calibration matches the historical equity premium. In contrast, EU risk preferences require a risk aversion coefficient of 34 to match the equity premium based on the same endowment, because they result in considerably less covariation between the pricing kernel and returns when consumption and dividends are imperfectly correlated. Within the representative agent paradigm, my study is the first to rationalize the equity premium without either large consumption shocks (such as disasters or long-run risks) or high risk aversion (as in the external habit model).

In the last part of the paper, I augment the IID model with long-run risks in the spirit of Bansal and Yaron (2004). The extended model replicates the amount of excess volatility and

return predictability in the data, while remaining consistent with the prices, returns, and Sharpe ratios of options with different strikes.

My study mainly relates to literatures on option prices in equilibrium models, on the measurement and modeling of risks in macroeconomic fundamentals, and on generalized disappointment aversion.

The main difficulty in rationalizing option *prices* in an equilibrium setting is that consumption and dividend data only display a moderate amount of left tail risk. Martin (2017) shows that even extensions of the Bansal and Yaron (2004) and Barro (2006) models that were explicitly designed to capture features of stock market volatility remain inconsistent with option prices. Specifically, he quantifies the amount of tail risk in the option-implied (risk-neutral) return distribution by the difference between the VIX and SVIX indices, and shows that the long-run risk models of Bollerslev et al. (2009) and Drechsler and Yaron (2010) *undershoot* this metric by an order of magnitude. In contrast, the variable disaster probability model of Wachter (2013) *overshoots* it by an order of magnitude.

The main difficulty in rationalizing option *returns* in an equilibrium setting is that standard utility functions rely on a single risk aversion parameter. As a result, it is not possible to increase investors' aversion against large tail events without simultaneously increasing their aversion against small tail events. But precisely such a change would be necessary to replicate the large negative returns of derivatives that insure against market corrections without overshooting the equity premium. Dew-Becker et al. (2017) show that the models of Drechsler and Yaron (2010) and Wachter (2013) both imply that Sharpe ratios of variance swaps (which can be interpreted as option portfolios) are comparable to that of the market, whereas they are 4 times larger in the data. Beason and Schreindorfer (2019) show that states associated with stock market crashes contribute less than one-twentieth to the equity premium in the Wachter (2013) model, but for almost its entirety in the data. Intuitively, a model's (in)ability to match option returns therefore coincides with its (in)ability to capture sources of market risk premiums in the data. Beason and Schreindorfer show

that the extended habit model of Bekaert and Engstrom (2017) does somewhat better, implying that almost 4/10 of the equity premium is associated with said return states, but it does so by overstating their probability by over 100%. In contrast, the high price of tail risk implied by GDA allows my model to remain consistent with the above-mentioned metrics. As such, it is the first model to match option returns and Sharpe ratios in a consumption-based framework.

Empirically, I document that dividend growth (and returns) are more correlated with consumption growth conditional on low consumption growth realizations. Left tail events for equity holders are therefore likely to coincide with left tail events in consumption. This fact is consistent with the finding of Berger et al. (2019), who show that times of high realized market volatility tend to coincide with low consumption growth. My results also relate to Duffee (2005), who shows that the conditional correlations between consumption growth and returns moves procyclically, and to Xu (2018), who shows that the conditional correlations between consumption and dividend growth moves procyclically. There are two important differences. First, downside correlations condition on an *ex post* event rather than on *ex ante* information. As such, they are easier to measure because they do not rely on a model to capture state dependence. Second, higher downside correlations make equity riskier, all else equal, whereas the channel highlighted by Duffee and Xu has the opposite implication.

Routledge and Zin (2010) show that recursive utility with GDA risk preferences gives rise to time-varying effective risk aversion in settings with persistent state variables. The subsequent literature has explored this channel quantitatively, showing that it results in realistic degrees of return predictability in endowment economies (Bonomo et al., 2011) as well as production economies (Liu and Miao, 2015). GDA also has been used to resolve household portfolio choice puzzles (Dahlquist et al., 2016; Andries and Haddad, 2019), and to provide a theoretical foundation for risk factors that explain cross-sectional differences in expected returns (Ang et al., 2006; Lettau et al., 2014; Delikouras, 2017; Farago and Tedongap, 2018; Delikouras and Kostakis, 2019). I contribute to the literature by illustrating that GDA can match the equity premium based on low

risk aversion, and by showing that GDA preferences can be calibrated to capture the large Sharpe ratios of derivatives that insure against market corrections, such as straddles, variance swaps, and out-of-the-money put options. Option *prices* were first analyzed in a model with GDA preferences by Schreindorfer (2014), but that study did not consider option *returns* or the implications of nonnormal shocks, and it is subsumed by the present work.

1 Macroeconomic Tail Risks

This section documents a nonlinearity in the relationship between firms' cash flows and macroeconomic shocks. I show that, because of this nonlinearity, modeling dividends as jointly normal with consumption significantly understates their exposure to macroeconomic tail events. A simple Gaussian-exponential mixture distribution is shown to capture the observed tail dependence.

1.1 Measuring tail risk exposure

To begin, define the downside correlation between two random variables X_t and Y_t as

$$\rho(x) = \text{corr}[X_t, Y_t | X_t \leq x], \tag{1}$$

which conditions on a left-tail event in X.[2] I use percentiles of X as thresholds. For example, $\hat{\rho}(p_{50})$ denotes the empirical correlation between X and Y in periods where X falls below its median.

I compute downside correlations between log consumption and dividend growth and between log consumption growth and log excess returns for annual U.S. data over 1930–2017. In both cases, consumption growth is used as the conditioning ("X") variable. To account for the well-known fact that returns move before fundamentals, I rely on a return time series that leads consumption

[2] Ang and Chen (2002) introduced downside correlations as a means to capture tail dependence between the returns on different assets. Different from their definition, (1) conditions on just one of the series (rather than both) falling below a given threshold in order to focus on tail events in consumption.

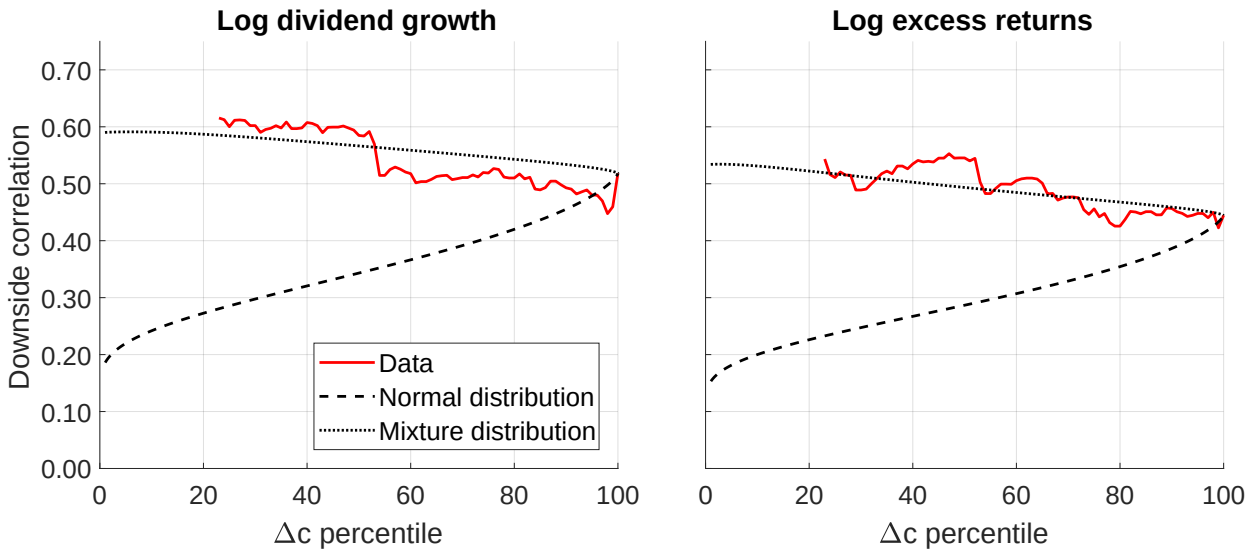

Figure I: Downside correlations

I plot downside correlations between log consumption and dividend growth (left panel) and between log consumption growth and log excess returns on the market (right panel). The horizontal axis shows the threshold for consumption below which the correlation is computed. Solid lines correspond to annual U.S. data over 1930–2017; dotted (dashed) lines correspond to normal-exponential mixture (normal) distributions that match the unconditional correlation of the two series. Empirical downside correlations are only computed above the 23rd percentile to ensure a minimum of 20 observations.

by two quarters, but measure dividend growth contemporaneously.[3] The solid lines in Figure I show that both dividends (left panel) and returns (right panel) are more strongly correlated with consumption in times of low consumption growth. The downside correlation for low consumption growth percentiles exceeds 0.5 in both cases, which implies that left tail events in consumption and dividends (and in consumption and returns) are likely to coincide.

The observation that correlations are higher in the left tail is important because the common assumption of joint normality implies the opposite pattern. The dashed lines in Figure I show downside correlations for a bivariate normal distribution that matches the unconditional corre-

[3]Backus et al. (2010) conduct a detailed analysis of the lead-lag relationship between returns and the business cycle and propose an approach for incorporating it into models.

lation in the data. For dividend growth, the correlation with consumption growth falls from an unconditional value of 0.52 to a downside correlation of 0.35 below the median and a value of less than 0.2 in the far left tail. The right panel shows a similar pattern for a normal distribution that is calibrated to consumption growth and returns. Because economic theory implies that left tail events in consumption are associated with elevated marginal utility, asset pricing models based on bivariate normal shocks imply that shocks to dividends carry unrealistically low risk premiums. Similarly, tail events in dividends and returns contribute unrealistically little to the equity premium.

1.2 Robustness and statistical significance

To illustrate the pervasiveness of the decreasing pattern in downside correlations, Table I compares unconditional correlations to downside correlations below the median for a variety of alternative time series. I consider payout (cash dividends plus repurchases), reinvested dividends, earnings, the post-war sample, a biannual data frequency, and dividends and returns for the five Fama-French industries. Additionally, macroeconomic conditions are measured by either consumption of gross domestic product (GDP) growth. Table I shows that, in 39 of 40 considered cases, the downside correlation exceed the corresponding unconditional correlation. Hence, the observation that equity is more correlated with macroeconomic fundamentals in bad economic times is a robust feature of the data.

To assess how precisely the decreasing pattern in downside correlations is measured, I compute the bootstrapped sampling distribution of $\hat{\rho}(p_{100}) - \hat{\rho}(p_{50})$ for the data in Figure I. The sampling distribution implies a 95% confidence interval of $[-0.235, 0.214]$ for dividends and $[-0.279, 0.214]$ for returns. To save space, I show it in Figure 1 of the Online Appendix.

The most efficient way to test the null of a particular distribution is to ask whether the data could have been generated by the distribution, that is, to impose the null. I do so via a Monte Carlo simulation that computes the finite sample distribution of $\rho(p_{100}) - \rho(p_{50})$ under the null of joint

Table I: Downside correlations for alternative time series – corr$[X_t, Y_t | X_t \leq x]$

Description	Asset	Y	X = Δc		X = Δy	
			$x = p_{50}$	$x = p_{100}$	$x = p_{50}$	$x = p_{100}$
Benchmark	MKT	Dividend growth	58.5	51.8	56.1	40.5
		Returns	54.5	44.5	46.4	35.3
		Payout growth	60.0	55.7	54.4	41.4
		Reinvested 1	66.3	61.5	56.9	45.7
		Reinvested 2	28.7	11.2	25.8	10.4
		Earnings growth	29.7	27.2	35.1	29.9
Post-war	MKT	Dividend growth	17.4	8.5	34.1	24.6
		Returns	59.4	34.2	57.6	42.4
Bi-annual	MKT	Dividend growth	61.3	54.0	56.3	37.3
		Returns	62.4	52.7	49.8	41.4
FF5 Industries	CNSMR	Dividend growth	38.5	39.9	38.1	25.5
		Returns	51.3	38.6	43.4	30.7
	MANUF	Dividend growth	53.0	50.1	49.9	40.0
		Returns	55.0	42.6	45.6	34.1
	HITEC	Dividend growth	9.3	3.7	22.4	5.1
		Returns	49.7	41.8	42.8	32.6
	HLTH	Dividend growth	32.6	17.0	0.2	-17.8
		Returns	46.7	32.3	29.1	14.5
	OTHER	Dividend growth	68.7	63.1	62.1	53.4
		Returns	56.0	44.6	48.8	37.5

Downside correlations below the median ("$x = p_{50}$" columns) are compared to unconditional correlations ("$x = p_{100}$" columns). The conditioning variable X is either log consumption (real nondurable and services consumption per capita) or output (real gross domestic product per capita) growth; the Y variable (column 3) is either a measure of log cash flow growth or the excess log return. Besides "Post-war" (annual, 1947–2017) and "Biannual" (biannual, 1930–2017), all series are annual and span 1930–2017. Dividend growth in year t is the sum of twelve monthly dividends in year t divided by the sum of twelve monthly dividends in year $t-1$; Payout equals cash dividends plus repurchases, as in Bansal et al. (2005). "Reinvested 1" ("Reinvested 2") is dividend growth, where monthly dividends are reinvested at the risk-free rate (the monthly market return) until the end of the year, like in Chen (2009); earnings come from Robert Shiller's Web site.

normality. When the normal distribution is calibrated to the unconditional correlation between consumption and dividend growth, the empirical difference of $\hat{\rho}(p_{100}) - \hat{\rho}(p_{50}) = -0.067$ has a p-value of 2.7%. For a calibration that matches the unconditional correlation between consumption growth and returns, the equivalent p-value is only 0.57%. When the same test is repeated for each of the 40 cases in Table I, the null hypothesis of joint normality is rejected at the 10% significance level in 34/40 cases, at the 5% level in 28/40 cases, and at the 1% level in 13/40 cases. These results are tabulated in the Online Appendix. Clearly, the normal distribution understates the likelihood of joint tail events in consumption and dividend growth, as well as the likelihood of joint tail events in consumption and returns.

1.3 Modeling tail risk exposure

Intuitively, two variables are more correlated on the downside if negative shocks are more likely to be common than positive shocks. Similarly, downside correlations are increasing in the left tail if shocks that are more negative are also more likely to be common. These properties can be captured with the Gaussian-exponential mixture distribution (henceforth "mixture distribution")[4]

$$\begin{aligned} \varepsilon^c &= \sqrt{1-\omega^2}\eta^c + \omega(\eta^e - 1) \\ \varepsilon^d &= \sqrt{1-\omega^2}\eta^d + \omega(\eta^e - 1), \end{aligned} \tag{2}$$

where η^c and η^d are standard normal, η^e is exponentially distributed with a unit rate parameter, and all three shocks are mutually independent. The distribution relies on a single free parameter, the mixture parameter ω. It is straightforward to show that ε^c and ε^d have a mean of zero, a standard deviation of one, a skewness of $2\omega^3$, and an unconditional correlation of ω^2. For $\omega < 0$, skewness is negative, the unconditional correlation is positive, and downside correlations are higher in the left tail because large negative values for ε^c and ε^d are more likely to be observed when the common exponential shock takes on a large positive value.

[4]Dahlquist et al. (2016) applied the same mixture model in a portfolio choice setting.

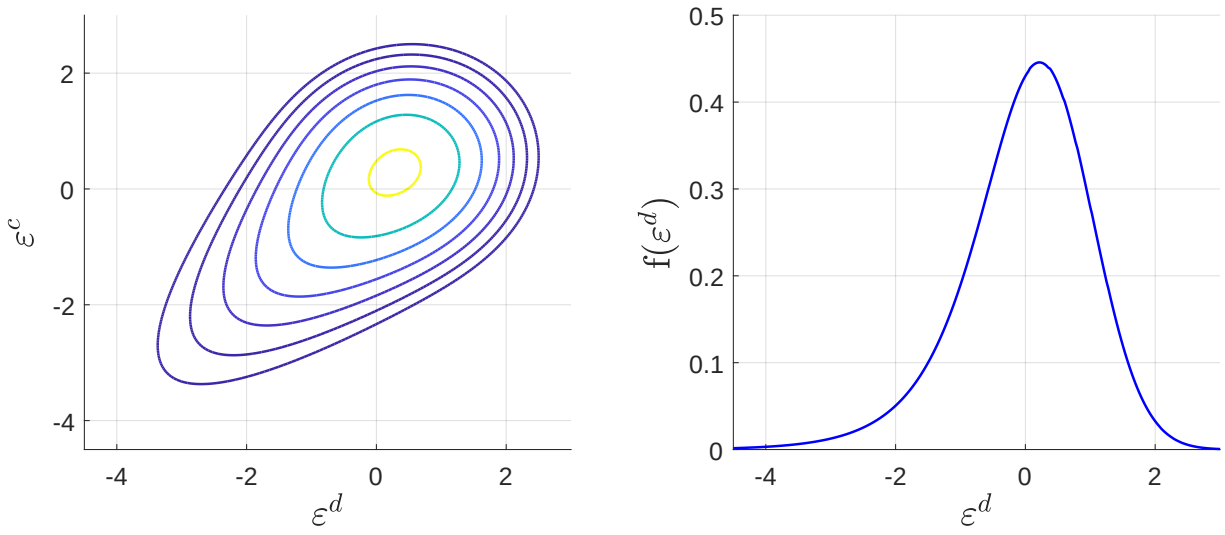

Figure II: Mixture distribution

The left panel shows a contour plot of the joint probability density function (PDF) of $(\varepsilon^c, \varepsilon^d)$ based on the mixture distribution in Equation (2) and $\omega = -0.73$. The right panel shows the corresponding marginal PDF, which has an identical shape for both variables.

Figure II illustrates the shape of the mixture distribution when it is calibrated to match the unconditional correlation between annual log consumption and dividend growth over 1930–2017. The contour plot of the joint distribution (left panel) is "avocado shaped," whereas the marginal distribution (right panel) is left-skewed. Both shapes contrast strongly with those of a bivariate normal distribution, which has elliptical contour lines and symmetric marginal distributions. The dotted line in the left panel of Figure I shows that the calibrated distribution does a nice job of capturing the empirically observed downside correlations between consumption and dividend growth. When the mixture distribution is instead calibrated to match the unconditional correlation between log excess returns and (lagged) log consumption growth, it is equally successful at capturing the downside correlations between the two series, as shown in the right panel of Figure I.

Similar to the analysis for the normal distribution, I test the null of the mixture distribution with a Monte Carlo simulation. When the mixture distribution is calibrated to the unconditional

correlation between consumption and dividend growth, the finite sample distribution of $\rho(p_{100}) - \rho(p_{50})$ implies a p-value of 42.1% for the empirical difference of $\hat{\rho}(p_{100}) - \hat{\rho}(p_{50}) = -0.067$. For returns, the equivalent p-value is 16.1%. When the same test is repeated for each case in Table I, the null hypothesis is rejected at the 10% significance level in 2/39 cases, at the 5% level in 1/39 cases, and never at the 1% level.[5] Hence, the mixture distribution successfully captures the joint tail behavior of consumption and dividend growth (and of consumption growth and excess returns) in the data. It also generates marginal distributions that are unimodal and left-skewed, as in the data. The next section proposes an asset pricing model based on shocks from the mixture distribution.

2 IID Model

I consider a representative-agent pure-exchange economy with a single nonstorable consumption good. This section details the model and discusses how it is solved. Analytical asset pricing formulae are derived in the appendix.

2.1 Assumptions

Aggregate consumption in period t equals C_t, whereas the aggregate dividend is D_t. Equity is a claim to the dividends in all future periods. Consumption and dividends follow random walks,

$$\begin{aligned} \Delta c_{t+1} \equiv \ln\left(\frac{C_{t+1}}{C_t}\right) &= g + \sigma\varepsilon^c_{t+1} \\ \Delta d_{t+1} \equiv \ln\left(\frac{D_{t+1}}{D_t}\right) &= g + \varphi\sigma\varepsilon^d_{t+1}. \end{aligned} \tag{3}$$

Both log growth rates have a constant mean of g and constant volatility. Consumption volatility equals $\sigma \geq 0$, whereas dividend volatility is higher and equal to $\varphi\sigma$. The parameter $\varphi \geq 1$ captures

[5]For dividend growth in the health care industry and GDP growth, the unconditional correlation is negative. Because the mixture distribution does not accommodate negative correlations, this test is omitted.

leverage, as in Abel (1999) and Campbell and Cochrane (1999). The innovations η^c and η^d are IID over time and given by the Gaussian-exponential mixture distribution in (2).

Following Epstein and Zin (1989), the representative agent's time t utility, V_t, is given by the constant elasticity of substitution (CES) recursion

$$V_t^\rho = (1-\beta)C_t^\rho + \beta\mu_t(V_{t+1})^\rho. \tag{4}$$

$0 < \beta < 1$ characterizes impatience and $\rho \leq 1$ measures the preference for intertemporal substitution (the elasticity of intertemporal substitution for deterministic consumption paths is $\frac{1}{1-\rho}$). Risk preferences are captured by the function $\mu_t(V_{t+1})$, which equals the certainty equivalent of random future utility using the time t conditional probability distribution. Henceforth, I use the shorthand notation $\mu_t \equiv \mu_t(V_{t+1})$. The certainty equivalent features Generalized Disappointment Aversion (GDA), as in Routledge and Zin (2010), and it is defined by the implicit function

$$u(\mu_t) = E_t\Big[u(V_{t+1})\Big] - \theta E_t\Big[\Big(u(\delta\mu_t) - u(V_{t+1})\Big)\mathbb{D}_{t+1}\Big], \tag{5}$$

where $\mathbb{D}_{t+1} \equiv \mathbf{1}\{V_{t+1} \leq \delta\mu_t\}$ is a disappointment indicator and the period utility function equals

$$u(x) = \begin{cases} \frac{x^\alpha - 1}{\alpha} & \text{for } \alpha \leq 1, \alpha \neq 0 \\ \ln(x) & \text{for } \alpha = 0 \end{cases}. \tag{6}$$

Equation (5) nests two well-known preference specifications as special cases. First, for $\theta = 0$ the second term drops out and risk preferences simplify to expected utility. In this case, the certainty equivalent is given by the *explicit* function $\mu_t = (E_t[V_{t+1}^\alpha])^{\frac{1}{\alpha}}$, the utility function equals the specification used in Epstein and Zin (1991) and Bansal and Yaron (2004), and $\alpha \leq 1$ captures static risk aversion (the coefficient of relative risk aversion for static gambles is $1-\alpha$). Second, for $\delta = 1$ risk preferences simplify to Gul's (1991) model of disappointment aversion (DA), which was developed to resolve the Allais (1979) paradox. The asset pricing implications of recursive utility with DA risk preferences have been analyzed by Epstein and Zin (2001) in an endowment economy

and by Campanale et al. (2010) in a production economy, but neither paper analyzes option prices or considers the implications of nonnormal shocks. DA preferences imply that all outcomes below the certainty equivalent are considered disappointing and receive a penalty, the magnitude of which is governed by the parameter $\theta > 0$. GDA preferences, which represent the most general version of Equation (5) ($\theta > 0$ and $\delta \neq 1$), place the disappointment threshold further into the tail of the conditional distribution of V_{t+1}. In particular, only realizations of V_{t+1} that fall below a fraction δ of the certainty equivalent μ_t are considered disappointing.

Because the risk aggregator in (5) and (6) is linearly homogeneous, CES+EU, CES+DA, and CES+GDA all fall in the class of Epstein and Zin (1989) preferences. To distinguish between the nested cases, the rest of the paper uses the terminology "EU," "DA," and "GDA" to refer to the risk preference component of the intertemporal utility function, with the understanding that time preferences are assumed to be of the CES type. Importantly, however, the terminology "EU" is not meant to indicate that the entire utility function is of the expected utility type, which is only true when time preferences are time-separable (for $\alpha = \rho$).

2.2 Solving the model

In equilibrium, the agent maximizes utility and the price of the consumption claim clears markets. The unit root in consumption makes utility nonstationary, so the model solution is characterized in terms of utility ratios rather than levels. Define the time-invariant ratios $v \equiv \frac{V_t}{C_t}$ and $m \equiv \frac{\mu_t(V_{t+1})}{C_t}$. The value-to-consumption ratio follows by dividing (4) by C_t,

$$v^\rho = (1-\beta) + \beta m^\rho. \tag{7}$$

Similarly, dividing (5) by C_t and re-arranging terms results in

$$m^\alpha = \frac{\mathbb{E}\left[v^\alpha e^{\alpha \Delta c_{t+1}}(1+\theta \mathbb{D}_{t+1})\right]}{1+\theta\delta^\alpha \mathbb{E}\left[\mathbb{D}_{t+1}\right]}. \tag{8}$$

The disappointment event $V_{t+1} \leq \delta\mu_t$ can be expressed in terms of consumption growth as

$$\Delta c_{t+1} \leq \ln\left(\frac{\delta m}{v}\right) \equiv x_v, \tag{9}$$

where x_v is a constant that depends on the equilibrium value of v (m can be substituted out based on Equation (7)). The model solution takes three steps. First, the expectations in (8) are evaluated analytically by integrating over the Gaussian and exponential innovations in consumption growth. Next, m is substituted out of the resulting expression based on (7), which yields a single nonlinear equation in v. I relegate the lengthy equation to (A.10) in the appendix. Lastly, the equation is solved numerically with a standard nonlinear equation solver. As discussed on page 1308 of Routledge and Zin (2010), the implicit function theorem guarantees that the nonlinear equation has a unique solution. The steps for the limiting case $\alpha \to 0$ are analogous and shown in Appendix A.

2.3 Asset Pricing

Determining asset prices requires the pricing kernel, which equals the representative agent's marginal rate of intertemporal substitution (IMRS), $M_{t+1} \equiv \frac{\partial V_t / \partial C_{t+1}}{\partial V_t / \partial C_t}$. Differentiating (4) with respect to C_t and C_{t+1}, respectively, and substituting back into M_{t+1} yields $M_{t+1} = \beta \left(\frac{v}{m}\right)^{1-\rho} \frac{\partial \mu_t}{\partial V_{t+1}}$. Using the implicit function theorem to differentiate μ_t in (5) with respect to V_{t+1} gives $\frac{\partial \mu_t}{\partial V_{t+1}} = \left(\frac{v}{m}\right)^{\alpha-1} e^{(\alpha-1)\Delta c_{t+1}} \frac{1+\theta\mathbb{D}_{t+1}}{1+\theta\delta^\alpha \mathbb{E}[\mathbb{D}_{t+1}]}$, so that the IMRS is given by

$$M_{t+1} = \beta_v \times \underbrace{e^{(\alpha-1)\Delta c_{t+1}}}_{\text{Small risks}} \times \underbrace{(1 + \theta \mathbf{1}\{\Delta c_{t+1} \leq x_v\})}_{\text{Tail risks}}, \tag{10}$$

where $\beta_v \equiv \beta \frac{(v/m)^{\alpha-\rho}}{1+\theta\delta^\alpha \mathbb{E}[\mathbb{D}_{t+1}]}$ is a constant that depends on the equilibrium value of v. The pricing kernel takes a very simple form, consisting of a constant, a "small risks" factor that is familiar from the EU case, and a "tail risks" factor that overweights left tail outcomes of consumption growth by a factor of $(1 + \theta)$. Note that the IID environment implies that the pricing kernel is a function

of consumption growth only. The benchmark calibration sets $\alpha = 1$, so that the small risks term drops out and the pricing kernel reduces further to a step function (see Figure IV). Given a solution for v, asset prices can be computed based on the Euler equation

$$E[M_{t+1}R^i_{t+1}] = 1, \tag{11}$$

where R^i denotes the gross return of asset i. Using this approach, Appendix A derives analytical expressions for the prices of risk-free bonds, dividend strips, equity, forwards, options with different strike prices, the VIX index, and the SVIX index of Martin (2017).

It is worth pointing out that the "long-run risks" term of the Epstein and Zin (1989) pricing kernel, $(v/m)^{\alpha-\rho}$, has been absorbed into the constant β_v due to the IID setting, rather than due to GDA risk preferences. In non-IID settings, this term is state-dependent and hence time-varying. For a discrete state, such as the one considered in Section 5, the value function can be found by solving a system of nonlinear equations, and asset prices continue to be computable analytically thereafter. In contrast, continuous-valued (affine or nonaffine) state processes do not allow for analytical asset pricing formulae, because the certainty equivalent equals an implicit function under GDA. Therefore, the value function and asset prices have to be computed numerically, either on a grid or via projection methods.

3 Quantitative Implications

This section evaluates the model's quantitative implications for growth rates, basic asset prices, and option market moments in a calibration.

3.1 Data

Consumption is measured by the sum of real nondurables and services consumption per capita from the Bureau of Economic Analysis (BEA), whereas nominal dividends are imputed from cum- and

ex-dividend market returns from the Center for Research in Security Prices (CRSP). Dividends and returns are converted to real terms using the personal consumption expenditure deflator from the BEA.[6] Following Beeler and Campbell (2012), the ex ante annual real risk-free rate is approximated by the difference between the yield of a 3-month Treasury bill and an inflation forecast from a predictive regression.[7]

Option price quotes from January 5, 1990 to December 29, 2017, are obtained from the Chicago Board Options Exchange (CBOE). The data set contains end-of-day information of all European exercise style options written on the S&P 500 (underlying SPX). I apply standard filters to remove observations with low liquidity and obvious data errors, and measure option prices by the midquote. Because the model is calibrated at a monthly frequency, option prices are interpolated to a constant maturity of 30 calendar days for each day in the sample. Both the filters and interpolation method are detailed in the Online Appendix. Option returns are computed on each day based on interpolated 30-day option prices and the return of the S&P 500 over the subsequent 30 calendar days. The unconditional moments reported below are based on this daily sample of overlapping 30-day option prices and returns.

3.2 Calibration and basic annual moments

The agent's decision interval is assumed to be monthly, and I target moments of annual (time-aggregated) quantity and return data over 1930–2017, and moments of 1-month options over 1990–2017. Dividend growth in year t is the sum of twelve monthly dividends in year t divided by the sum of twelve monthly dividends in year $t-1$, and similarly for consumption. Annual log returns equal the sum of twelve monthly log returns, both for the market and for risk-free bonds. To compute

[6] Using the consumer price index as a deflator leads to a mean dividend growth rate significantly lower than the mean growth rate of consumption, whereas the PCE deflator avoids this issue. I thank an anonymous referee for pointing this out.

[7] See the appendix to Beeler and Campbell (2012) for details.

Table II: Calibration

Fundamentals				Preferences			
$g \times 12$	$\sigma \times \sqrt{12}$	φ	ω	β^{12}	ρ	θ	δ
0.0182	0.026	5.22	-0.73	0.985	$1 - 1/1.5$	9.76	0.975

The table shows parameter values for the model's monthly calibration. The period utility function is $u(x) = x$, i.e. the calibration assumes $\alpha = 1$ in (6).

moments in the model, I simulate one million samples of equivalent length as the data, compute the moment of interest in each sample, and then report the median value.

Table II summarizes the calibration, and panels A and B of Table III shows moments for annual growth rates and asset prices. I set $g = 0.0182/12$ to match the average consumption growth rate and assume the same growth rate for dividends. The volatility of consumption and dividend growth are matched by setting $\sigma = 0.026/\sqrt{12}$ and $\varphi = 5.22$, while their unconditional correlation is replicated based on a mixture parameter of $\omega = -0.73$. The four endowment parameters are thus uniquely pinned down by four quantity moments.

Time aggregation of an IID random walk results in positive first-order correlation (Working 1960). The model calibration results in an autocorrelation of 0.23 for both consumption and dividend growth, close to the empirical autocorrelation of 0.19 for dividends. Consumption growth has a higher autocorrelation in the data, but Kroencke (2017) points out out that this is a result of data filtering in the National Income and Product Accounts, and shows that unfiltered consumption data has about the same autocorrelation as dividends. Overall, the random walk model thus provides a very good and parsimonious description of the basic quantity moments in the data.

The rate of time preference of $\beta = 0.985^{1/12}$ together with a high EIS of $1/(1-\rho) = 1.5$ results in a low annual risk-free rate of about 1%. Hence, the model resolves the risk-free rate puzzle via the well-known channel of Weil (1989). Different from models with long-run risks, however, the EIS value is inconsequential for risk premiums because of the IID environment. In particular, for

Table III: Moments

Panel A: Annual growth rates

	$\mathbb{E}[\Delta c]$	$\sigma[\Delta c]$	$ac1[\Delta c]$	$\mathbb{E}[\Delta d]$	$\sigma[\Delta d]$	$ac1[\Delta d]$	$corr[\Delta c, \Delta d]$
Data	1.82	2.11	0.50	1.71	11.02	0.19	0.53
Model	1.82	2.11	0.23	1.82	11.02	0.23	0.53

Panel B: Annual asset prices

	$\mathbb{E}[R^f]$	$\mathbb{E}[R - R^f]$	$\sigma[R - R^f]$	$\mathbb{E}[pd]$
Data	0.46	8.02	19.75	3.42
Model	1.04	8.02	14.59	2.83

Panel C: Higher moments of monthly market returns

	$\overline{skew_t^*[R^{ex}]}$	$\overline{kurt_t^*[R^{ex}]}$	$skew[R^{ex}]$	$kurt[R^{ex}]$	VIX-SVIX
Data	−1.22	6.21	−0.88	5.09	0.57
Model	−1.16	5.31	−0.72	4.19	0.44

Panel D: Monthly straddle returns

	$\mathbb{E}[R^s]$	$\sigma[R^s]$	$skew[R^s]$	$kurt[R^s]$	$SR[R^s]$
Data	−187.30	217.86	1.30	6.12	−0.86
Model	−180.30	236.41	1.37	5.61	−0.77

Moments of annual log growth rates expressed as a percentage (panel A). Moments of annual returns expressed as a percentage and of the annual log price-dividend ratio (panel B). Unconditional physical moments and average conditional risk-neutral moments of monthly ex-dividend market returns (panel C). Annualized return moments for 1-month ATM straddles (panel D). Annual data moments for panels A and B are based on the 1930–2017 sample, whereas monthly data moments for panels C and D are based on the 1990–2017 option sample. Model moments are small sample medians.

any value of ρ, the pricing kernel in (10) is observationally equivalent to one with a different value of ρ and a compensating adjustment of β and δ.

I set $\alpha = 1$ so that risk preferences have no curvature, the only source of risk aversion consists of disappointments, and the pricing kernel in (10) reduces to a step function. This assumption eliminates one free parameter, and it simplifies the intuition behind the utility function, but it is inconsequential for the main results. In particular, moderate degrees of curvature have a negligible effect on asset pricing moments, similar to their effect in the well-known EU case (Mehra and Prescott 1985). The disappointment magnitude ($\theta = 9.76$) and threshold ($\delta = 0.975$) are set to jointly match the equity premium, and to match the average returns of put options with different strikes as well as possible. Specifically, it is possible to lower the disappointment threshold and simultaneously increase the disappointment magnitude to keep the equity premium unchanged. Doing so increases the average returns of put options with low strikes relative to those with high strikes because it increases the convexity of the pricing kernel as a function of returns, a mechanism I discuss in Section 4.1. The risk preference calibration results in a disappointment probability of 0.54% per month, so that disappointment events occur once every $\frac{1}{12\times 0.0054} = 15.4$ years on average.

The IID setting results in a transparent model mechanism and, as we will see next, it successfully replicates many option moments. Because it implies constant valuation ratios, however, the IID model does not generate excess volatility or return predictability: The volatility of annual returns is 14.59% (vs. 19.75% in the data). In Section 5, I show that a simple extension makes the model consistent with return predictability while preserving its consistency with option market data.

3.3 Option prices

The payoff of a 1-period put option with strike K equals $S_t \times \max\{0, X - R^{ex}_{t+1}\}$, where $X = K/S_t$ denotes moneyness and $R^{ex}_{t+1} = S_{t+1}/S_t$ is the ex-dividend market return. Plugging the payoff

function into the Euler equation (11) shows that the put price,

$$P_t(X) = S_t \times \mathbb{E}_t[M_{t+1} \max\{0, X - R^{ex}_{t+1}\}], \tag{12}$$

reflects the joint distribution of returns and the pricing kernel in the left tail of the market return distribution. Puts are more valuable when (a) returns below X are more likely (so that expected put cash flows are higher) or when (b) returns below X are more likely to coincide with high values of M (so that put risk premiums are more negative). While option prices do not allow us to gauge the relative importance of both channels, prices for different strikes contain a wealth of information about the joint distributions of (R^{ex}, M). Because this joint distribution also determines the equity premium, a plausible theory of stock market risk premiums should be consistent with option prices.

3.3.1 Implied volatilities

I express option prices in units of annualized Black and Scholes (1973) implied volatilities (IVs), which makes lognormality a useful reference point. In particular, models in which (R^{ex}, M) are conditionally jointly lognormal, such as Campbell and Cochrane (1999) or Bansal and Yaron (2004), satisfy the assumptions of Black and Scholes (1973) over the agent's decision interval and therefore result in IVs that do not vary with moneyness. The left panel of Figure III shows that, in contrast to this case, the GDA-mixture model captures the steep negative slope of the empirical IV curve (the "IV skew") almost perfectly. The difference between the IVs of an out-of-the-money (OTM) put with moneyness $X = 0.9$ and an at-the-money (ATM) put with moneyness $X = 1$ equals 8.6 percentage points in data and 8.4 percentage points in the model. The level of the IV curve, however, is slightly lower than in the data due to the model's lack of excess volatility.[8]

[8]Similar to Drechsler and Yaron (2010) and other studies that examine option moments in equilibrium settings, I calibrate the model to fundamentals over 1930–2017, a period over which the volatility of monthly market returns was about 3 percentage points higher than over the 1990–2017 option sample. Because a higher volatility translates to higher option prices, this effect counterbalances the model's lack of excess volatility somewhat. For the same reason, however, the *level* of the IV curve (or of the VIX) is arguably not a suitable metric for evaluating models whose fundamentals are calibrated to a longer sample.

Figure III: Equity index options. Average implied volatilities (left; annualized percentage), average returns (middle; percentage), and Sharpe ratios (right; annualized) for 1-month put options are plotted against moneyness. Put returns are given by $\max\{0, K - S_{t+1}\}/P_t(K) - 1$ and moneyness by K/S_t. Model moments represent small sample medians based on 1 million samples of equivalent length as the 1990–2017 option sample.

3.3.2 Risk-neutral moments

For comparison with other studies, it is useful to note that the IV curve reflects the same information as the risk-neutral return distribution. In particular, Breeden and Litzenberger (1978) show that the second derivative of the put price with respect to moneyness, multiplied by $S_t \times R_t^f$, equals the conditional risk-neutral probability density function (PDF) of the ex-dividend market return:

$$f_t^*(R^{ex}) = S_t \times R_t^f \times \frac{\partial^2 P_t(X)}{\partial X^2}\bigg|_{X=R^{ex}}.$$

At any point in time, the relative prices of options with different strikes (the IV curve) therefore contain the same information as the time-t conditional moments of the risk-neutral PDF. Because the model is able to replicate the shape of the average IV curve, it is therefore not surprising that it also provides a good match for risk-neutral moments. Panel C of Table III shows that, on average, the conditional risk-neutral distribution has a large negative skewness of $\overline{skew_t^*[R^{ex}]} = -1.16$ (vs. -1.22 in the data) and a kurtosis of $\overline{kurt_t^*[R^{ex}]} = 5.31$ (vs. 6.21 in the data). To evaluate whether these moments are based on a plausible mechanism, one would ideally check whether the model is

also consistent with higher moments of the *conditional* physical return distribution. Because these moments cannot be computed in the data without additional assumptions (a model), I instead rely on *unconditional* moments. Ex-dividend returns in the model have a skewness of -0.72 and a kurtosis of 4.19, quite close to the empirical values of -0.88 and 5.09.

VIX minus SVIX. Martin (2017) shows that the degree of nonnormality of the risk-neutral distribution can alternatively be summarized by the difference between the VIX index and a related simple VIX (or "SVIX") index. The squared VIX equals the conditional entropy of the risk-neutral return distribution and the squared SVIX equals its conditional variance, so that $VIX - SVIX$ reflects a (nonlinear) weighted average of higher risk-neutral moments. Martin uses $VIX - SVIX$ to evaluate whether the models of Campbell and Cochrane (1999), Bansal and Yaron (2004), Bansal et al. (2012), Bollerslev et al. (2009), Drechsler and Yaron (2010), and Wachter (2013) can capture the shape risk-neutral distribution in the data, and concludes that "none of [the six models] come close to matching [...] the mean level of VIX minus SVIX observed in the data." The author's table IV shows that the average $VIX - SVIX$ is an order of magnitude too large in the Wachter (2013) model, and an order of magnitude too small in the remaining models. In contrast, panel C of Table III shows that the GDA-mixture model provides a good match for Martin's diagnostic, implying an average value of 0.44 (vs. 0.57 in the 1990–2017 sample).

Option prices, summarized in terms of either average implied volatilities, conditional risk-neutral moments, or the VIX-SVIX diagnostic, show that the risk-neutral return distribution is decisively negatively skewed. By themselves, however, they do not tell us whether the risk-neutral skewness results from conditional skewness in the physical return distribution (high expected put cash flows) or from high marginal utility in states of low market returns (low put risk premiums).

3.4 Option returns

The average returns of options with different strikes are useful for disentangling the effects of expected cash flows (physical tail risk) and risk premiums (risk preferences) on option prices. For

intuition on the information content of option returns, it is helpful to consider the example of an ATM put option and an OTM put option. The ATM put derives most of its expected cash flows from fairly likely, small, negative market returns, and only a small fraction from fairly unlikely, large, negative market returns. In contrast, the OTM put (with, say, $X = 0.9$) derives all its expected cash flows from fairly unlikely, large, negative market returns. If OTM puts earn substantially more negative average returns than ATM puts, we can therefore conclude that, on average, large negative market returns must coincide with higher marginal utility than small negative market returns. Hence, the average returns of puts with different strikes reveal how the pricing kernel varies with the level of market returns. In turn, an understanding of this covariation lies at the heart of the equity premium puzzle.

The middle panel of Figure III shows that average put returns are negative, large in magnitude, and considerably lower (more negative) for low moneyness levels. The model almost perfectly replicates this pattern. For example, puts with a moneyness of $X = 1$ have average *monthly* returns of -31.0% in the model (vs. -30.9% in the data), whereas puts that are 6% OTM have average returns of -69.4% in the model (vs. -71.4% in the data).[9] The right panel of Figure III shows that Sharpe ratios on put options behave very similar to returns, being negative, large in magnitude, and considerably lower for low moneyness levels. Once again, the model successfully replicates the empirical pattern. ATM puts have an annualized Sharpe ratio of -0.83 (vs. -0.81 in the data), whereas puts that are 6% OTM have a Sharpe ratio of -1.36 (vs. -1.44 in the data). For perspective, it is helpful to note that the market earned a much smaller Sharpe ratio of 0.54 over the 1990–2017 option sample.

[9]For comparison, Broadie et al. (2009, table II) find very similar average returns of -28.9% and -61.9% for puts with the same moneyness levels over August 1987–June 2006.

3.4.1 Variance risk premium

Many recent studies examine risk premiums on option portfolios that hedge against market volatility. The best-known example is an ATM straddle, which consists of one put and one call option with a moneyness of $X = \frac{K}{F_{t,T}}$ (where $F_{t,T}$ is the forward price for maturity T) and a common maturity of T. Panel D of Table III shows that the model successfully replicates the entire distribution of 1-month straddle returns, which have an annualized mean of -176% (vs. 187% in the data), an annualized volatility of 237% (vs. 218% in the data), a skewness 1.37 (vs. 1.30 in the data), a kurtosis of 5.60 (vs. 6.12 in the data), and an annualized Sharpe ratio of -0.75 (vs. -0.86 in the data).[10]

To interpret these numbers, note that the payoff of an ATM straddle, $\max\{0, K - S_{t+1}\} + \max\{0, S_{t+1} - K\} = S_t \times |R^{ex}_{t+1} - 1|$, increases in the realized volatility of the market. Because volatility is persistent, a common interpretation is that straddle returns reflect investors' aversion against shocks to expected market volatility (e.g., Coval and Shumway 2001). This interpretation is intuitive, because increased volatility reflects worse investment opportunities going forward. Recent evidence by Dew-Becker et al. (2017), however, shows that option portfolios that are only exposed to innovations in *expected* market volatility and immune to *realized* market volatility have historically earned average returns and Sharpe ratios close to zero. As noted by the authors, this implies that investors are averse against high *realized* volatility, but not against high *expected* volatility. Because large positive returns do not coincide with high marginal utility, an even simpler interpretation is that investors worry about large drops in the market. Arguably, the "variance risk premium" is therefore better described as a "crash risk premium." In line with this interpretation, Bollerslev

[10]The variance premium is sometimes defined as the difference between the squared VIX (divided by 12) and the conditional variance of monthly returns. In my calibration, this difference equals 9.04/10,000, which is close to the empirical value of 10.55/10,000 over 1990–2009 reported by Drechsler (2013). Relative to this metric, straddles have the advantage of representing an implementable trading strategy without assuming that stock prices cannot jump.

and Todorov (2011) estimate that 88.4% of the variance premium is associated with returns of -10% or less over horizons of a few weeks. The evidence from straddles and put options shows that the model in this paper is quantitatively consistent with this crash risk premium in options.

4 Inspecting the Mechanism

I illustrate the implications of alternative risk preferences and shock specifications with a series of calibrations and assess how the level of risk aversion under GDA compares to that of more common utility functions.

4.1 Risk preferences

To understand the role of GDA risk preferences, it is helpful to consider the model's pricing kernel. The left panel of Figure IV shows the pricing kernel as a function of log consumption growth, along with a consumption growth distribution that is scaled to fit into the same figure. Because the pricing kernel equals a step function that only takes on elevated values in the far left tail, risk premiums in the model exclusively reflect exposure to macroeconomic tail risks (left tail events in Δc). For perspective, it is helpful to compare the GDA pricing kernel to the one implied by standard EU risk preferences. The dashed line shows the EU pricing kernel for a calibration with identical endowment and time preference parameters, and a risk aversion coefficient of $1-\alpha = 34.25$ that is chosen to match the equity premium. The EU pricing kernel increases monotonically for low values of consumption growth but, despite the high risk aversion coefficient, aversion against tail events is considerably lower than under GDA. Macroeconomic tail risks are therefore less important for risk premiums than under GDA, while small drops in consumption are more important.

To understand risk premiums on the dividend claim in the model, it is helpful to assess how marginal utility varies with the level of market returns. I do so by examining the "projected pricing kernel," $\mathbb{E}[M|R^{ex}]$. The law of iterated expectations implies that the projected kernel has the same

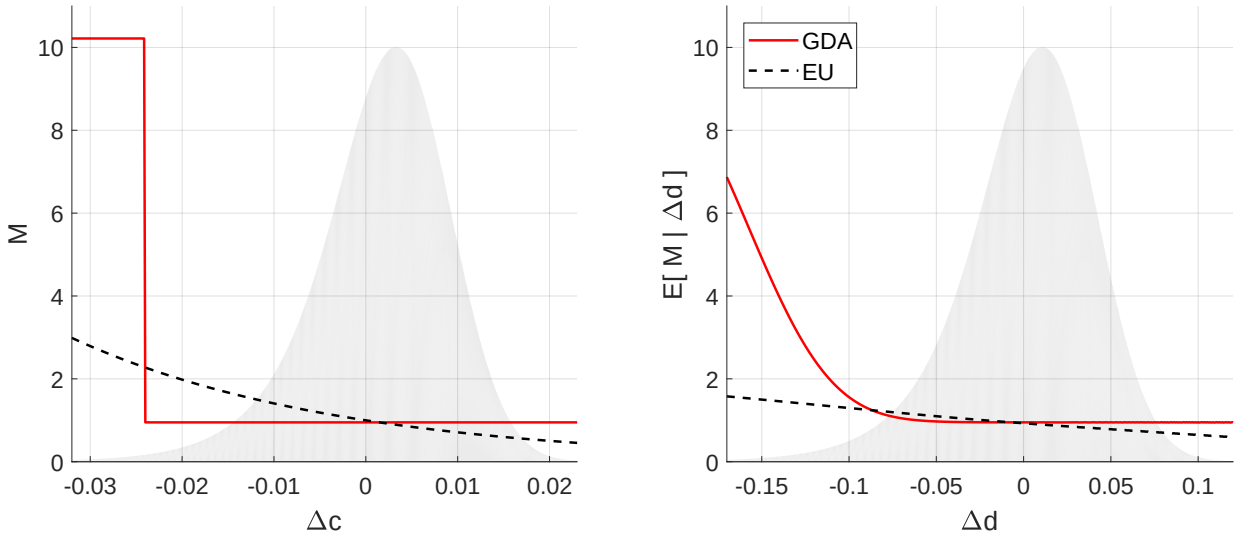

Figure IV: Pricing kernel

The left panel shows the pricing kernel as a function of log consumption growth, and the right panel shows the conditional mean of the pricing kernel as a function of log dividend growth. Solid lines correspond to the calibration in Table II. Dashed lines correspond to a calibration with expected utility risk preferences ($\theta = 0$ and $\alpha = -33.25$) and otherwise identical parameters that also matches the equity premium. Both panels overlay the corresponding densities, scaled to fit in the same graph.

pricing implications as M for assets with payoffs that solely depend on R^{ex}, that is,

$$\mathbb{E}\left[Mg(R^{ex})\right] = \mathbb{E}\left[\mathbb{E}\left[M|R^{ex}\right]g(R^{ex})\right] \tag{13}$$

for any payoff function $g(\cdot)$. Because index options are written on the ex-dividend price index, put risk premiums for different strikes are informative about the shape of the projected pricing kernel.

The right panel of Figure IV shows the projected pricing kernel as a function of log ex-dividend returns, $\ln R^{ex} = \Delta d$, and reveals a number of important facts about the model mechanism. First, imperfect correlation between Δc and Δd implies that $\mathbb{E}[M|R^{ex}]$ is a continuous, smooth, and monotonically decreasing function of returns, because the conditional disappointment probability changes smoothly with returns. This is a desirable model property because empirical estimates of $\mathbb{E}[M|R^{ex}]$ show that, rather than resembling a step function, the pricing kernel is approximately

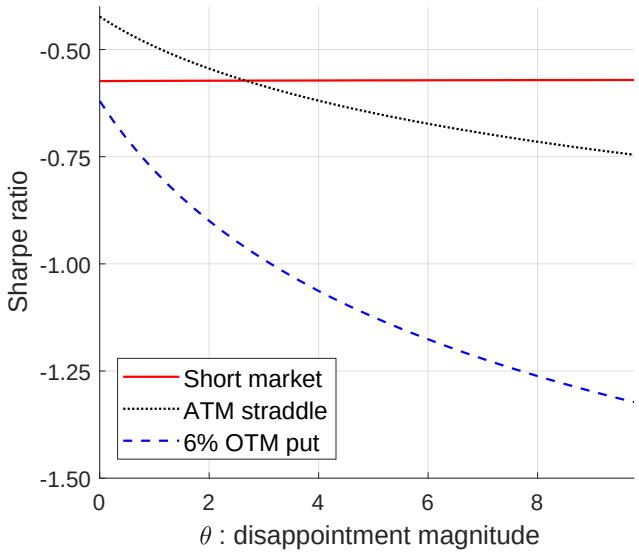

Figure V: Risk preferences and the price of risk

The figure shows annualized Sharpe ratios of different assets for alternative risk preference calibrations. I start with the calibration in Table II, lower θ, and adjust α to hold the equity premium fixed. EU risk preferences (the benchmark GDA calibration) correspond to the leftmost (rightmost) point on the horizontal axis.

flat for positive and moderately negative returns and steeply downward sloping in the far left tail (see, e.g., Ait-Sahalia and Lo (2000, Figure 3) or Rosenberg and Engle (2002, figure 6)). Second, relative to EU, the GDA kernel takes on slightly lower values for small negative market returns and considerably higher values for large negative returns. GDA therefore implies that a larger part of the equity premium is associated with states of very low market returns and that risk premiums and Sharpe ratios of put options decrease more quickly for lower moneyness levels than under EU.

I conduct a simple experiment ("Experiment 1") to illustrate explicitly how risk preferences affect the Sharpe ratios of various assets. Starting with the benchmark calibration in Table II ($\theta = 9.76$ and $\alpha = 1$), I gradually lower the disappointment magnitude (θ) and simultaneously adjust the curvature parameter (α) to hold the equity premium fixed. Parameters for the endowment and

time preferences are kept unchanged, so that all considered calibrations have realistic implications for quantity growth rates, the risk-free rate, and the equity premium. Figure V shows annualized Sharpe ratios of (a short position in) the market, an ATM straddle, and a 6% OTM put option as a function of the disappointment magnitude.[11] The benchmark GDA calibration (the rightmost point on the horizontal axis) results in Sharpe ratios of -0.57 (short market), -0.75 (ATM straddle), and -1.32 (OTM put), close to the empirical values of -0.54, -0.86, and -1.44 over 1990–2017. Lowering the disappointment magnitude keeps the market Sharpe ratio unchanged and moves the Sharpe ratios of straddles and put options closer to that of the market. For EU risk preferences (the leftmost point on the horizontal axis), the three assets have very similar Sharpe ratios of -0.57 (short market), -0.42 (ATM straddle), and -0.62 (OTM put). Contrary to the data, EU risk preferences therefore imply that Sharpe ratios of straddles are smaller than that of the market!

Dew-Becker et al. (2017) provide similar evidence for *non-IID* models with EU risk preferences. In particular, the models of Drechsler and Yaron (2010) and Wachter (2013) both imply that variance swaps (which can be interpreted as option portfolios) earn Sharpe ratios comparable to that of the market, whereas they are 4 times larger in the data. Because option Sharpe ratios reflect the joint distribution of market returns and the pricing kernel, the findings presented here and by Dew-Becker et al. call into question whether EU risk preferences are consistent with sources of market risk premiums in the data.[12] Related evidence comes from Beason and Schreindorfer (2019), who show based on options data that 85/100 of the equity premium is associated with states that feature stock market crashes, while such states contribute one-tenth or less in the models of Bansal and Yaron (2004), Barro (2006), and Wachter (2013). Intuitively, the issue with EU risk preferences lies

[11]Note that the calibrations differ in both θ and α. The Online Appendix shows the values of α that correspond to each value of θ in Figure V.

[12]Dew-Becker et al. (2017) show that a calibration of the disaster model with time-varying recovery rates in Gabaix (2012) is consistent with the high Sharpe ratios of option portfolios, despite relying on EU risk preferences. Because Dew-Becker et al. almost double the RRA coefficient relative to the original calibration, however, Gabaix's model is also not *jointly* consistent with Sharpe ratios of options and the market.

in the fact that they make it impossible to increase the relative importance of large shocks for risk premiums. The above evidence shows that the GDA model in the present paper is consistent with option Sharpe ratios, while Beason and Schreindorfer show that it is consistent with their equity premium decomposition (said states contribute 79/100 to the equity premium in the model).

4.2 Shock distributions

The negative slope of the implied volatility curve shows that returns are left-skewed under the risk-neutral measure. So far, however, it is unclear whether the model replicates this fact due to large physical tail risks (mixture shocks) or due to a large price of tail risk (GDA risk preferences).

The effect of physical tail risks (left-skewness and high downside correlations) in the model can be illustrated by increasing the mixture parameter ω, which pushes the joint distribution of log consumption growth and returns closer to a bivariate normal distribution. The drawback of such an experiment, however, is that higher values of ω are also associated with a lower unconditional correlation and therefore a lower equity premium, all else equal. To overcome this effect, I now assume that Gaussian innovations to consumption and dividend growth have a correlation of $\varrho \equiv \text{Corr}[\eta^c, \eta^d]$, and I adjust risk preference parameters to hold the equity premium fixed. Starting at the benchmark calibration in Table II ($\omega = -0.73$, $\varrho = 0$), the experiment ("Experiment 2")...

- increases the mixture parameter ω to decrease the amount of tail risk in innovations
- increases ϱ to hold the unconditional correlation between annual consumption and dividend growth fixed at the empirical value of 0.53.
- increases the disappointment magnitude θ to hold the equity premium fixed at the empirical value of 8.02% per year
- increases the disappointment threshold δ to hold the disappointment probability fixed at 0.54% per month – the value in the benchmark calibration.

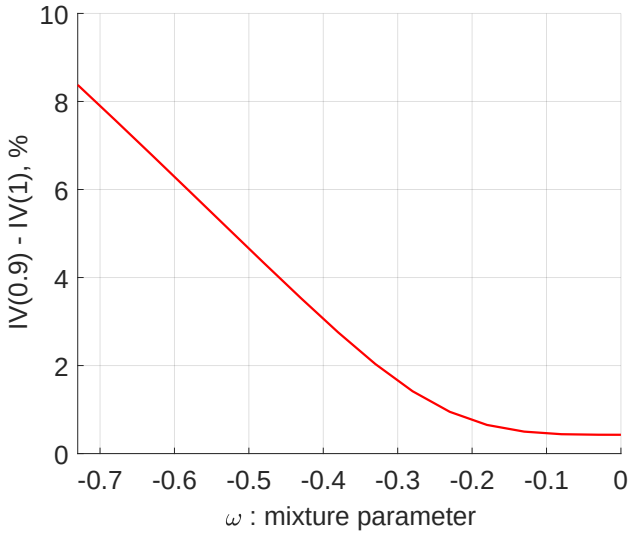

Figure VI: Shock distributions and the implied volatility skew

The difference between implied volatilities at moneyness 0.9 and moneyness 1 is plotted for different shock calibrations. I start with the calibration in Table II, increase ω, and simultaneously (a) increase ϱ to hold the correlation between consumption and dividend growth fixed, (b) increase θ to hold the equity premium fixed, and (c) increase δ to hold the disappointment probability fixed. Normal shocks (the benchmark mixture-shock calibration) correspond to the rightmost (leftmost) point on the horizontal axis.

The remaining four parameters in Table II are left unchanged. Note that, if risk preference parameters were not adjusted in the experiment, calibrations with higher values of ω would also imply lower risk premiums despite identical unconditional correlations because (a) less skewness makes disappointment events rarer and (b) lower downside correlations make disappointments less likely to coincide with large drops in the market. In contrast, the adjustment of (θ, δ) ensures that all considered calibrations have similar implications for the equity premium, the risk-free rate, and the first two moments of quantity growth rates.

The red line in Figure VI shows the implied volatility skew—measured by the difference in IVs at moneyness 0.9 and moneyness 1—as a function of ω.[13] The benchmark calibration (the

[13]Note that the calibrations differ in terms of ϱ, θ, δ, and ω. The Online Appendix shows the values of

leftmost point on the horizontal axis) results in an IV skew of 8.4%, while an increase in ω results in a monotonic decrease in the skew. For bivariate normal shocks (the rightmost point on the horizontal axis), the skew is barely positive at 0.4%, despite GDA preferences. Relative to the well-known case with normal shocks and EU risk preferences, which implies a perfectly horizontal IV curve (a skew of zero), adding GDA preferences therefore has a negligible effect on the skew. The intuitive reason for this result is that the probability of left tail returns quickly approaches zero under normality, so that deep OTM put options have very small expected cash flows. The high price of risk associated with low market returns under GDA does not suffice to offset this small probability. Clearly, the model requires nonnormal shocks to match option prices, while Section 4.1 showed that it requires GDA risk preferences to match option Sharpe ratios.

4.3 Risk aversion

Routledge and Zin (2010) show that effective risk aversion can be time-varying under GDA, but the previous literature has not compared the *level* of risk aversion under GDA to that of more common utility functions. This section provides such a comparison and shows that (a) GDA requires less risk aversion than EU risk preferences to match the equity premium and (b) the benchmark GDA calibration implies the same level of risk aversion as an EU calibration with a relative risk aversion coefficient of about ten.

It is important to note that dynamic settings do not allow us to quantify risk aversion *for an individual utility function*, but rather only to make relative comparisons of risk aversion *across different utility functions*. Specifically, while the Arrow-Pratt measures of risk aversion allow for a separation between risk and risk aversion in static settings, the same is not true for dynamic environments. The reason lies in how utility is measured. In static settings, utility is quantified by an atemporal utility function $u(C)$ over consumption that reflects preferences only, and relative risk aversion (RRA) can be approximated by the local curvature $-C\frac{u_{cc}}{u_c}$. In dynamic settings,

ϱ, θ, and δ that correspond to each value of ω in Figure VI.

utility is quantified by an intertemporal value function $V(W)$ over wealth that reflects preferences, technology, and market structures, and the local curvature of the value function, $-W\frac{V_{ww}}{V_w}$, does not isolate risk aversion (see, e.g., the discussion in Campbell and Cochrane, 1999).

My comparison of risk aversion across utility functions is based on the welfare costs of risk, which I define as the fraction of consumption that the agent would give up (today and at every future date and state) in order to exchange her endowment for an alternative endowment with the same mean consumption growth rate but no risk:[14]

$$w = 1 - \frac{v}{v^{no-risk}}. \tag{14}$$

$v^{no-risk}$ is the value-to-consumption ratio for the risk-free endowment.[15] To isolate risk aversion, I consider calibrations that assume identical time preference and endowment parameters and have risk preferences that match the equity premium. Specifically, starting with the benchmark calibration in Table II, I lower the disappointment magnitude (θ) and adjusting the curvature parameter (α) to keep the equity premium fixed (like in Experiment 1 in Section 4.1).

The left panel of Figure VII shows the welfare costs in (14) as a function of the disappointment magnitude. The benchmark GDA calibration ($\alpha = 1, \theta = 9.76$), which corresponds to the rightmost point on the horizontal axis, implies that the agent would give up 29.68% of her endowment to eliminate its riskiness. Lowering θ leads to a monotonic increase in welfare costs, which reach 68.07% for the nested case with EU risk preferences ($\alpha = -33.25, \theta = 0$).[16] Because the amount of

[14]Here, welfare costs measure the same object as "risk premiums" in Epstein et al. (2014). Lucas (1987) uses $\frac{w}{1-w}$ to measure the benefits of eliminating risk rather than w to measure its cost.

[15]$v^{no-risk}$ can be computed analytically: iterating on (4) yields $v^{no-risk} = \left(\frac{1-\beta}{1-\beta\mathbb{E}[e^{\Delta c}]^\rho}\right)^{1/\rho}$. Using Lemma 2.B.i in the appendix, expected consumption growth equals $\mathbb{E}[e^{\Delta c}] = \frac{\exp\{g-\sigma\omega+0.5\sigma^2(1-\omega^2)\}}{1-\sigma\omega}$.

[16]While these welfare costs are orders of magnitude larger than the welfare costs of business cycles in Lucas's study, they are similar to those in other endowment asset pricing models. In particular, Epstein et al. (2014) compute welfare costs of 57% for the model of (Bansal and Yaron 2004, case II), 29% for Barro (2009), and 65% for Wachter (2013). The reason lies in the fact that asset pricing studies assume higher risk aversion (and/or more endowment risk) than Lucas in order to match the equity premium.

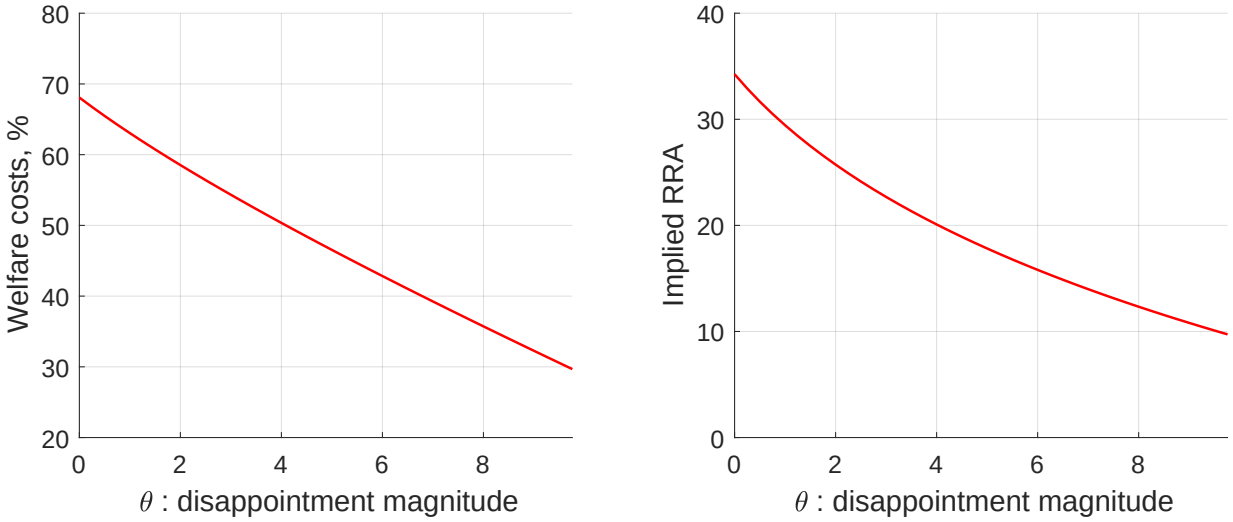

Figure VII: Measures of risk aversion

Left panel: starting with the calibration in Table II, I lower θ and adjust α to keep the equity premium fixed. The plot shows the welfare costs in (14) for each resulting calibration as a function of θ. Right panel: starting with the welfare costs implied by a GDA calibration in the left panel, I solve for the value of α that results in the same welfare costs under EU. The plot shows the implied level of RRA $(1-\alpha)$ as a function of θ. Both panels: EU risk preferences (the benchmark GDA calibration) correspond to the leftmost (rightmost) point on the horizontal axis.

endowment risk is identical for all considered calibrations, the comparison shows that risk aversion is monotonically decreasing in θ.

To understand the intuition behind this result, note that welfare costs are a function of *consumption risk*, while the equity premium depends on the covariation between *dividend risk* and the pricing kernel. Under imperfect correlation, shocks that result in elevated values for the GDA pricing kernel (infrequent, large, negative consumption shocks) are also very likely to coincide with adverse outcomes for dividends, especially for mixture shocks. It therefore requires relatively little risk aversion to produce a sizable equity premium. For EU, many shocks that result in elevated values for the pricing kernel (frequent, small, negative consumption shocks) have a sizable chance of coinciding with positive dividend growth, so it requires more risk aversion to match the equity

premium.

Because risk aversion under EU has been extensively discussed in existing work, it is useful to compare the level of effective risk aversion under GDA to a level of RRA $(1-\alpha)$ under EU. To do so, I start with the welfare costs implied by one of the GDA calibrations considered above, for example $w(\alpha = 1, \theta = 9.76) = 29.68\%$ in the benchmark calibration. Holding time preference and endowment parameters fixed, I then solve for the value of α that results in identical welfare costs under EU, for example, $w(\alpha = -8.72, \theta = 0) = 29.68\%$. Because the GDA and EU calibrations are based on an identical amount of endowment risk and result in identical welfare costs, they also imply the same degree of risk aversion. The GDA agent in the benchmark calibration therefore has the same degree of risk aversion as an EU agent with a RRA of $1-(-8.72) = 9.72$. The right panel of Figure VII shows the implied coefficient of RRA for each GDA calibration in the left panel. In line with the intuition that welfare costs in the experiment only change as a result of differences in risk aversion, implied RRA decreases monotonically with θ.

To evaluate how sensitive these findings are to the assumed endowment process, I repeat the analysis for an endowment with normal shocks in consumption and dividends and the same unconditional correlation. I find that, when GDA preferences are calibrated to match the equity premium and the disappointment probability of the calibration in Table II (as in Experiment 2 in Section 4.2), they imply welfare costs of 30.65% and the same risk aversion as EU risk preferences with $1-\alpha = 10.29$. The measurement of implied risk aversion under GDA is therefore not significantly affected by the assumed shock structure.

5 Dynamic Model

The combination of GDA risk preferences and mixture shocks allows for an explanation of the market risk premium that is consistent with options data. Because the IID model implies constant valuation ratios, however, it generates neither excess volatility nor return predictability. In this

section, I show that a simple dynamic extension of the model can quantitatively replicate these data features, while remaining consistent with the prices and returns of the market and of equity index options.

5.1 Model and calibration

I augment the model with time-variation in macroeconomic risk by assuming that log consumption and dividend growth follow

$$\begin{aligned} \Delta c_{t+1} =& g + x_t + \sigma_t \varepsilon^c_{t+1} \\ \Delta d_{t+1} =& g + \phi x_t + \varphi \sigma_t \varepsilon^d_{t+1}, \end{aligned} \tag{15}$$

where η^c and η^d are IID over time and given by the mixture distribution (2). As in Bansal and Yaron (2004), x_t and σ_t capture persistent variation in the conditional mean and volatility of growth rates, respectively, and ϕ represents a source of leverage. An increase in ϕ makes dividend growth more volatile, more persistent, and more correlated with consumption growth. To keep the model solution tractable, I follow Bhamra et al. (2010) in assuming that time-variation in x_t and σ_t is driven by a two-state Markov chain,

$$x_t \in \begin{bmatrix} \chi \\ -\chi \end{bmatrix} \qquad \sigma_t \in \begin{bmatrix} -\xi \\ \xi \end{bmatrix} + \sigma \qquad P = \begin{bmatrix} \pi & 1-\pi \\ 1-\pi & \pi \end{bmatrix}. \tag{16}$$

The parameters $\chi \geq 0$ and $\xi \geq 0$ govern the amount of time-variation in conditional moments, whereas π governs their persistence. The "good" state (state 1) is characterized high expected growth and low uncertainty and the "bad" state (state 2) is characterized by low expected growth and high uncertainty. In this setting, the value-to-consumption ratio v_t can take on two (state-dependent) values, which follow from the solution of two nonlinear equations (the analogues of (A.10) for both states). Asset prices continue to be available analytically. In particular, expectations in all asset pricing formulae can be evaluated by first conditioning on the current and future

Table IV: Calibration of the dynamic model

Fundamentals				Preferences			
$g \times 12$	$\sigma \times \sqrt{12}$	φ	ω	β^{12}	ρ	θ	δ
0.0182	0.0214	5.9	-0.72	0.985	$1 - 1/1.5$	3.3	0.985
$\chi \times 12$	$\xi \times \sqrt{12}$	ϕ	π				
0.013	0.0071	2.2	0.999				

The table shows parameter values for the model's monthly calibration. The period utility function is $u(x) = x$; that is, the calibration assumes $\alpha = 1$ in Equation (6).

state, using results from the IID model to integrate over $(\varepsilon_c, \varepsilon_d)$, and then evaluating the remaining expectation over future states as a matrix product.

Table IV shows the model's monthly calibration. As in the IID case, $(g, \sigma, \varphi, \omega)$ are chosen to match average consumption growth, the volatility of consumption and dividend growth, and the correlation between both series. Because x_t (together with ϕ) affects the correlation between consumption and dividends as well as their volatilities, however, parameter values differ slightly different from those in the IID model: I set $g = 0.0182/12$, $\sigma = 0.00214/\sqrt{12}$, $\varphi = 5.9$, and $\omega = -0.72$. Time preferences are kept unchanged relative to the IID case ($\beta = 0.985^{1/12}$ and $\rho = 1 - 1/1.5$), while risk preferences are chosen to match the equity premium and to match the average returns of put options with different strikes as well as possible ($\delta = 0.985$ and $\theta = 3.3$). The persistence of the Markov chain is set to replicate the first-order autocorrelation of the log price-dividend ratio ($\pi = 0.999$) and the parameters controlling the values of x_t and σ_t are chosen to match the volatilities of the log price-dividend ratio and excess returns ($\chi = 0.013/12$, $\xi = 0.0071/\sqrt{12}$, $\phi = 2.2$). The calibration implies a conditional disappointment probability of 0.3% in state 1 (one event every 27.6 years) and 4.5% in state 2 (one event every 1.9 years).

5.2 Results

Panel A of Table V shows the model's implications for annual growth rates. As before, all model moments represent small sample medians. By construction, the calibration provides a good match for the first two moments of consumption and dividend growth, including their correlation. The addition of x_t makes time-aggregated growth rates more persistent than in the IID case, but less so than in typical calibrations of the long-run risks model due to a lower volatility of x_t. The first-order autocorrelation of annual log dividend growth equals 0.25 and that of consumption equals 0.34.

Panel B shows annual asset pricing moments. Like in the IID setting, the model provides a good match for the means of the risk-free rate, the equity premium, and the log price-dividend ratio. Additionally, the dynamic model replicates the low volatility of the risk-free rate, the high volatility and persistence of the log price-dividend ratio, and the "excess volatility" of stock returns. The volatility equals 1.53% for the risk-free rate (vs. 2.81% in the data), 0.48 for the log price-dividend ratio (vs. 0.46 in the data), and 19.88% for excess returns (vs. 19.75% in the data), while the autocorrelation of the log price-dividend ratio equals 0.88 (vs. 0.89 in the data). A large part of the excess return volatility results from time variation in effective risk aversion, which is low in state 1 because disappointments (high marginal utility events) are relatively unlikely, and high in state 2 because disappointments are relatively likely. Time-varying risk aversion maps into time-varying risk premiums, which makes the price dividend ratio (and hence returns) more volatile than in models with EU risk preferences, such as Bansal and Yaron (2004). This channel is discussed in detail in Routledge and Zin (2010).

Panel C evaluates the model's ability to capture the predictability of long-horizon returns. I regress log excess returns for horizons between 1 and 10 years on the log price-dividend ratio. In the model, slope coefficients decrease from -0.08 (vs. -0.07 in the data) at the 1-year horizon to -0.60 (vs. -0.58 in the data) at the 10-year horizon. A near-perfect match. Regression R^2's increase

Table V: Moments in the dynamic model

Panel A: Annual Growth Rates

	$\mathbb{E}[\Delta c]$	$\sigma[\Delta c]$	$ac1[\Delta c]$	$\mathbb{E}[\Delta d]$	$\sigma[\Delta d]$	$ac1[\Delta d]$	$corr[\Delta c, \Delta d]$
Data	1.82	2.11	0.50	1.71	11.02	0.19	0.53
Model	1.84	2.13	0.34	2.04	11.03	0.25	0.53

Panel B: Annual asset prices

	$\mathbb{E}[R^f]$	$\sigma[R^f]$	$\mathbb{E}[R - R^f]$	$\sigma[R - R^f]$	$\mathbb{E}[pd]$	$\sigma[pd]$	$ac1[pd]$
Data	0.46	2.81	8.02	19.75	3.42	0.46	0.89
Model	1.32	1.53	8.17	19.88	3.25	0.48	0.88

Panel C: Return predictability

Horizon in years		1	2	4	6	8	10
$\hat{\beta}$	Data	-0.07	-0.16	-0.27	-0.38	-0.50	-0.58
	Model	-0.08	-0.15	-0.28	-0.39	-0.50	-0.60
R^2	Data	0.03	0.08	0.16	0.23	0.30	0.31
	Model	0.02	0.04	0.08	0.11	0.14	0.15

Moments of annual log consumption and dividend growth expressed as a percentage (panel A). Moments of annual returns expressed as a percentage, and the annual log price-dividend ratio (panel B). Statistics from regressing multiperiod excess log returns on the lagged log price-dividend ratio, $\sum_{h=1}^{H}(r_{t+h} - r_{t+h}^f) = \alpha + \beta \times pd_t + \varepsilon_{t+h}$ (panel C). The data frequency is annual, and results for $H > 1$ are based on overlapping samples. All panels: The data span 1930–2017. Model moments are small sample medians.

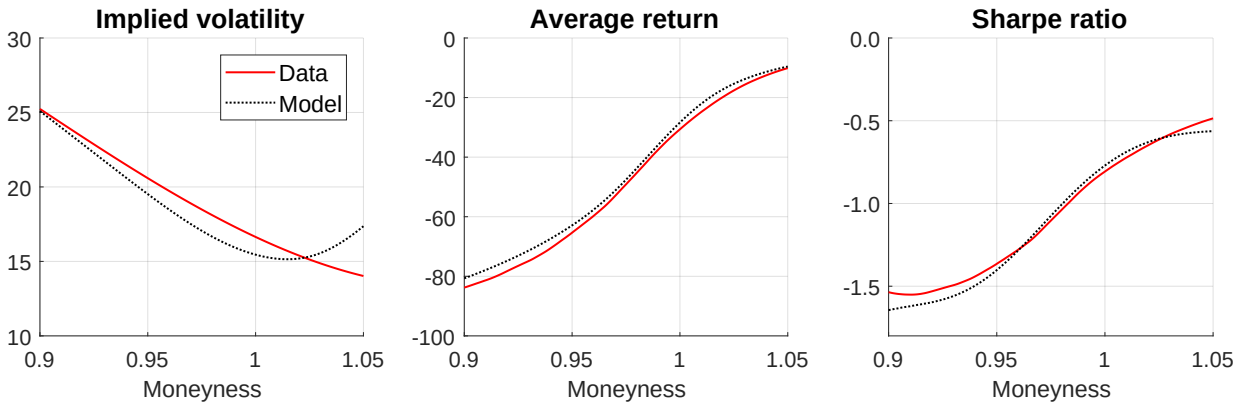

Figure VIII: Equity index options in the dynamic model

Average implied volatilities (left; annualized percentage), average returns (middle; percentage), and Sharpe ratios (right; annualized) for 1-month put options are plotted against moneyness. Put returns are given by $\max\{0, K - S_{t+1}\}/P_t(K) - 1$ and moneyness by K/S_t. Model moments represent small sample medians based on 1 million samples of equivalent length as the 1990–2017 option sample.

in the return horizon, but somewhat less so than in the data. The R^2 equals 2% (vs. 3% in the data) at the 1-year horizon and 15% (vs. 31% in the data) at the 10-year horizon.

Figure VIII evaluates option moments in the dynamic model. The model provides a good match for the shape of the average IV curve. The difference in implied volatilities of an OTM put with moneyness 0.9 and an ATM put with moneyness 1 equals 9.5 percentage points (vs. 8.6 percentage points in the data). One shortcoming of the model is that it results in a slight "smile" in IVs for moneyness levels above one, a pattern that is not present in the data. The middle and right panels show that the mean returns and Sharpe ratios of put options with different strikes are replicated almost perfectly. As in the IID model, large negative market returns therefore tend to coincide with much higher values of marginal utility than small negative returns. More broadly, the good match for average option returns suggests that the model is able to reproduce the degree of covariation between the pricing kernel and market returns in the data.

Overall, the results in this section show that the addition of time-varying macroeconomic risks

to the GDA-mixture model results in realistic amounts of excess volatility and return predictability, without affecting the model's consistency with options market data.

6 Conclusion

This paper presents a new consumption-based explanation for risk premiums in stock markets. In contrast to the Gaussian random walk model, but as in reality, I assume that consumption and dividend growth rates are highly correlated in bad states of the world. In contrast to the EU model of risk preferences, but as in reality, investors have asymmetric preferences over gains and losses. Together, the two assumptions result in a large equity premium because they imply that states of high marginal utility are very likely to coincide with bad news for equity holders.

The proposed mechanism differs from existing explanations of the equity premium along two key dimensions. First, it is consistent with a multitude of moments from index option markets, including the high Sharpe ratios of "crash insurance." In line with the empirical evidence of Bollerslev and Todorov (2011) and Beason and Schreindorfer (2019), a large part of the equity premium is therefore associated with states of low (but nondisastrous) market returns. Second, the model illustrates that it is possible to rationalize the equity premium without large consumption risks or high risk aversion. As such, it is the first representative agent model that respects both aspects of the Mehra and Prescott (1985) puzzle.

Many avenues suggest themselves for future work. Empirically, it is interesting to ask whether high downside correlations in fundamentals can explain the high downside correlations between equity returns documented by Ang and Chen (2002) and Longin and Solnik (2001). Theoretically, it is important to find a microfoundation for the tail dependence of fundamentals. In this spirit, Acemoglu et al. (2017) argue that the nonnormality of GDP growth is an important property of business cycles and show that it can be explained by the interplay of idiosyncratic shocks and sectoral heterogeneity in a general equilibrium model. Because their model abstracts from investment,

however, it cannot speak to consumption tail risk or the cotail risk between different macroeconomic series. A model that captures the joint tail behavior of different macro series could generate new insights about the origin of recessions.

Appendix. Solution details

A.1 Notation. Throughout the appendix, I denote the cumulative density function (cdf) of a univariate standard normal random variable by $\Phi(\cdot)$, and the corresponding probability density function (pdf) by $\Phi'(\cdot)$. The cdf of a bivariate standard normal vector with correlation r is denoted by $\Phi_r(\cdot,\cdot)$. Let $\mathbf{1}\{\cdot\}$, $sgn(\cdot)$, and $\overline{sgn}(\cdot)$ denote the indicator, sign (signum), and modified sign functions respectively:

$$\mathbf{1}\{x \leq a\} = \begin{cases} 1 & \text{if } x \leq a \\ 0 & \text{if } x > a \end{cases}$$

$$sgn(x) = \begin{cases} -1 & \text{if } x < 0 \\ 0 & \text{if } x = 0 \\ +1 & \text{if } x > 0 \end{cases} \qquad \overline{sgn}(x) = \begin{cases} -1 & \text{if } x \leq 0 \\ +1 & \text{if } x > 0 \end{cases}$$

For notational convenience, define

$$\kappa \equiv \frac{\omega}{\sqrt{1-\omega^2}} \tag{A.1}$$

$$\tilde{\sigma} \equiv \sigma\sqrt{1-\omega^2} \tag{A.2}$$

$$\beta_v \equiv \beta \frac{(v/m)^{\alpha-\rho}}{1+\theta\delta^{\alpha}\mathbb{E}[\mathbb{D}_{t+1}]} \tag{A.3}$$

$$x_v \equiv \ln\left(\frac{\delta m}{v}\right) = \ln(\delta) + \frac{1}{\rho}\left[\beta^{-1} + v^{-\rho}(1-\beta^{-1})\right] \tag{A.4}$$

$$z_v \equiv \frac{x_v - g}{\tilde{\sigma}} + \kappa \tag{A.5}$$

$$w_v \equiv sgn(\kappa)\exp\left\{\frac{1}{2\kappa^2} - \frac{z_v}{\kappa}\right\}\Phi\left(sgn(\kappa)z_v - \frac{1}{|\kappa|}\right) \tag{A.6}$$

A.2 Two Helpful Lemmas. Derivations heavily rely on the following lemmas, the proof of which can be found in the Online Appendix.

LEMMA 1 – TRUNCATED NORMALS.

Let x, y be standard normal with correlation ρ, and let a, b, c, d, r, s be constants. Then

(A) $\mathbb{E}\left[e^{rx+sy}\mathbf{1}\{a \le x \le b\}\mathbf{1}\{c \le y \le d\}\right]$

$$= e^{\frac{1}{2}(r^2+2\rho rs+s^2)}\left[\Phi_\rho(b^*, d^*) + \Phi_\rho(a^*, c^*) - \Phi_\rho(a^*, d^*) - \Phi_\rho(c^*, b^*)\right],$$

where $a^* = a - r - \rho s$, $b^* = b - r - \rho s$, $c^* = c - \rho r - s$, and $d^* = d - \rho r - s$.

(B) $\mathbb{E}\left[e^{rx}y^2\mathbf{1}\{x \le a\}\right] = e^{\frac{r^2}{2}}(1 + r^2\rho^2)\Phi(a - r) - e^{\frac{2ra-a^2}{2}}\frac{(r+a)\rho^2}{\sqrt{2\pi}}$

(C) $\mathbb{E}\left[e^{rx}y\mathbf{1}\{x \le a\}\right] = \rho e^{\frac{r^2}{2}}\left(r\Phi(a - r) - \frac{1}{\sqrt{2\pi}}e^{-\frac{(a-r)^2}{2}}\right)$

(D) $\mathbb{E}\left[x\mathbf{1}\{x \le a\}\right] = -\Phi'(a)$

LEMMA 2 – EXPONENTIAL RANDOM VARIABLES WITHIN A NORMAL DENSITY.
Let z be exponentially distributed with rate parameter 1, let a, b, c, d, f be constants with $a < 1$, and define $\lambda \equiv \frac{a-1-bc}{|c|}$ and $\lambda_f \equiv \frac{a-1-df}{|f|}$. Then

(A) $\mathbb{E}[e^{az}z\Phi(b+cz)]$
$$= \begin{cases} \frac{1}{(1-a)^2}\left(\Phi(b) + \frac{1-a}{c}\Phi'(b) + \left(\text{sgn}(c) + \frac{\lambda(1-a)}{c}\right)\exp\left\{\frac{\lambda^2}{2} - \frac{b^2}{2}\right\}\Phi(\lambda)\right) & , c \neq 0 \\ \frac{\Phi(b)}{(1-a)^2} & , c = 0 \end{cases}$$

(B) $\mathbb{E}[e^{az}\Phi(b+cz)] = \begin{cases} \frac{1}{1-a}\left(\Phi(b) + \text{sgn}(c)\exp\left\{\frac{\lambda^2}{2} - \frac{b^2}{2}\right\}\Phi(\lambda)\right) & , c \neq 0 \\ \frac{\Phi(b)}{1-a} & , c = 0 \end{cases}$

(C) $\mathbb{E}[\Phi'(b+cz)] = \frac{1}{|c|}\exp\left\{\frac{b}{c} + \frac{1}{2c^2}\right\}\Phi\left(-\text{sgn}(c)b - \frac{1}{|c|}\right)$, $c \neq 0$

(D) $\mathbb{E}[e^{-az^2+bz}] = e^{\frac{(b-1)^2}{4a}}\sqrt{\frac{\pi}{a}}\Phi\left(\frac{b-1}{\sqrt{2a}}\right)$

(E) $\mathbb{E}[e^{az}\Phi(b+cz, d+fz)] \equiv \Omega(a,b,c,d,f)$

$$= \frac{1}{1-a}\Bigg[\Phi_\rho(b,d) + sgn(c)\exp\left\{\frac{\lambda^2}{2} - \frac{b^2}{2}\right\}\left(\mathbf{1}\{\rho c > f\}\Phi(\beta_1) - \overline{sgn}(\rho c - f)\Phi_{r_1}(\gamma_1, \beta_1)\right) + sgn(f)\exp\left\{\frac{\lambda_f^2}{2} - \frac{d^2}{2}\right\}\left(\mathbf{1}\{\rho f > c\}\Phi(\beta_2) - \overline{sgn}(\rho f - c)\Phi_{r_2}(\gamma_2, \beta_2)\right)\Bigg]$$

where

$$\begin{array}{llll} r_1 = \frac{|\rho - \frac{f}{c}|}{\sqrt{\left(\rho - \frac{f}{c}\right)^2 + (1-\rho^2)}} & & r_2 = \frac{|\rho - \frac{c}{f}|}{\sqrt{\left(\rho - \frac{c}{f}\right)^2 + (1-\rho^2)}} & \\ \beta_1 = \frac{(1-a)\left(\rho - \frac{f}{c}\right) + cd - bf}{c\sqrt{\left(\rho - \frac{f}{c}\right)^2 + (1-\rho^2)}} & & \beta_2 = \frac{(1-a)\left(\rho - \frac{c}{f}\right) + fb - dc}{f\sqrt{\left(\rho - \frac{c}{f}\right)^2 + (1-\rho^2)}} & \\ \gamma_1 = \overline{sgn}(\rho c - f)\frac{bc + 1 - a}{|c|} & & \gamma_2 = \overline{sgn}(\rho f - c)\frac{df + 1 - a}{|f|} & \end{array}$$

Special cases that will prove useful below are:

(1.A.i) $\mathbb{E}[e^{rx+sy}\mathbf{1}\{x \leq a\}] = e^{\frac{1}{2}(r^2 + 2\rho rs + s^2)}\Phi(a - r - \rho s)$

(1.A.ii) $\mathbb{E}[e^{rx+sy}\mathbf{1}\{a \leq x\}] = e^{\frac{1}{2}(r^2 + 2\rho rs + s^2)}\left[1 - \Phi(a - r - \rho s)\right]$

(1.A.iii) $\mathbb{E}[e^{rx+sy}\mathbf{1}\{x \leq a\}\mathbf{1}\{y \leq b\}] = e^{\frac{1}{2}(r^2 + 2\rho rs + s^2)}\Phi(a - r - \rho s, b - \rho r - s)$

(1.C.i) $\mathbb{E}[e^{rx}y] = \rho r e^{\frac{r^2}{2}}$

A.3 Value Function. The disappointment event can be written as

$$V_{t+1} \leq \delta\mu_t \Leftrightarrow \Delta c_{t+1} \leq x_v \Leftrightarrow \eta^c_{t+1} \leq z_v - \kappa\eta^e_{t+1}. \tag{A.7}$$

Using first the definition of the standard normal cumulative distribution function and then Lemma 2.B, the disappointment probability is given by

$$\begin{aligned}\mathbb{E}[\mathbb{D}_{t+1}] =& \mathbb{E}\left[\mathbf{1}\left\{\eta^c_{t+1} \leq z_v - \kappa\eta^e_{t+1}\right\}\right] \\ =& \mathbb{E}\left[\Phi\left(z_v - \kappa\eta^e_{t+1}\right)\right] \\ =& \Phi\left(z_v\right) - \operatorname{sgn}(\kappa)\exp\left\{\frac{1}{2\kappa^2} - \frac{z_v}{\kappa}\right\}\Phi\left(\operatorname{sgn}(\kappa)z_v - \frac{1}{|\kappa|}\right) \\ =& \Phi\left(z_v\right) - w_v,\end{aligned} \tag{A.8}$$

where w_v is defined in (A.6). Integrating over the shocks in the numerator of (8) using first Lemma 1.A.i and next Lemma 2.B,

$$\begin{aligned}&\mathbb{E}[v^\alpha e^{\alpha\Delta c_{t+1}}(1+\theta\mathbb{D}_{t+1})] \\ &= e^{\alpha(g-\omega\sigma)+\frac{1}{2}\alpha^2\tilde{\sigma}^2}v^\alpha\mathbb{E}\left[e^{\alpha\omega\sigma\eta^e_{t+1}}(1+\theta\Phi(z_v-\kappa\eta^e_{t+1}-\alpha\tilde{\sigma}))\right] \\ &= e^{\alpha(g-\omega\sigma)+\frac{1}{2}\alpha^2\tilde{\sigma}^2}v^\alpha\frac{1+\theta\Phi\left(z_v-\alpha\tilde{\sigma}\right)-\theta\exp\left\{z_v\alpha\tilde{\sigma}-\frac{1}{2}\alpha^2\tilde{\sigma}^2\right\}w_v}{1-\alpha\omega\sigma}\end{aligned} \tag{A.9}$$

Finally, substituting (A.8) and (A.9) into (8), and then substituting out m from (7) yields a nonlinear equation in v,

$$1 = v^\alpha\frac{e^{\alpha(g-\omega\sigma)+\frac{1}{2}\alpha^2\tilde{\sigma}^2}\left[1+\theta\Phi\left(z_v-\alpha\tilde{\sigma}\right)-\theta\exp\left\{z_v\alpha\tilde{\sigma}-\frac{1}{2}\alpha^2\tilde{\sigma}^2\right\}w_v\right]}{\left[\beta^{-1}(v^\rho-1)+1\right]^{\frac{\alpha}{\rho}}\left(1-\alpha\omega\sigma\right)\left(1+\theta\delta^\alpha\Phi\left(z_v\right)-\theta\delta^\alpha w_v\right)}. \tag{A.10}$$

The equation can be solved with a standard nonlinear equation solver. Once v has been found, m follows from (7), which allows for the computation of the constants x_v and z_v that appear in the disappointment probability and the analytical asset pricing formulae that follow.

The limiting case $\alpha \to 0$. For $\alpha = 0$, the period utility function equals $u(x) = \ln(x)$, and similar to Equation (8), the scaled certainty equivalent is

$$\ln m = \frac{\mathbb{E}\left[\left(\ln v + \Delta c_{t+1}\right)\left(1+\theta\mathbb{D}_{t+1}\right)\right] - \theta\ln\delta\mathbb{E}[\mathbb{D}_{t+1}]}{1+\theta\mathbb{E}\left[\mathbb{D}_{t+1}\right]}. \tag{A.11}$$

To integrate over the cash flow shocks in the numerator of this expression, I begin by computing some intermediate objects. Using the definition of the normal cumulative distribution function and Lemma 1.D to integrate over η^c_{t+1},

$$\begin{aligned}\mathbb{E}[\Delta c_{t+1}\mathbb{D}_{t+1}] =& (g-\sigma\omega)\mathbb{E}\left[\mathbb{D}_{t+1}\right] + \tilde{\sigma}\mathbb{E}\left[\eta^c_{t+1}\mathbb{D}_{t+1}\right] + \sigma\omega\mathbb{E}\left[\eta^e_{t+1}\mathbb{D}_{t+1}\right] \\ =& (g-\sigma\omega)\mathbb{E}\left[\mathbb{D}_{t+1}\right] - \tilde{\sigma}\mathbb{E}\left[\Phi'(z_v-\kappa\eta^e_{t+1})\right] + \sigma\omega\mathbb{E}\left[\eta^e_{t+1}\Phi(z_v-\kappa\eta^e_{t+1})\right]\end{aligned}$$

Next, using Lemmas 2.A and 2.C to integrate over η^e_{t+1} and then re-arranging terms, taking into account (A.8),

$$\begin{aligned}
&\mathbb{E}[\Delta c_{t+1}\mathbb{D}_{t+1}] \\
&=(g-\sigma\omega)\mathbb{E}\left[\mathbb{D}_{t+1}\right]-\frac{\tilde{\sigma}}{|\kappa|}\mathrm{e}^{\frac{1}{2\kappa^2}-\frac{z_v}{\kappa}}\Phi\left(\mathrm{sgn}(\kappa)z_v-\frac{1}{|\kappa|}\right) \\
&\quad+\sigma\omega\left(\Phi(z_v)-\frac{1}{\kappa}\Phi'(z_v)-\mathrm{sgn}(\kappa)\left(1+\frac{z_v}{\kappa}-\frac{1}{\kappa^2}\right)\mathrm{e}^{\frac{1}{2\kappa^2}-\frac{z_v}{\kappa}}\Phi\left(\mathrm{sgn}(\kappa)z_v-\frac{1}{|\kappa|}\right)\right) \\
&=g\Phi(z_v)-\tilde{\sigma}\Phi'(z_v)+(\sigma\omega+x_v)\,\mathrm{e}^{\frac{1}{2\kappa^2}-\frac{z_v}{\kappa}}\Phi\left(\mathrm{sgn}(\kappa)z_v-\frac{1}{|\kappa|}\right)
\end{aligned} \tag{A.12}$$

From (A.12) and the expression for the conditional disappointment probability in (A.8), it follows that

$$\begin{aligned}
&\mathbb{E}[(\ln v+\Delta c_{t+1})\mathbb{D}_{t+1}] \\
&=(g+\ln v)\Phi(z_v)-\tilde{\sigma}\Phi'(z_v)-(\sigma\omega+\ln(\delta m))\,w_v,
\end{aligned} \tag{A.13}$$

so that the numerator of (A.11) can be written as

$$\begin{aligned}
&\mathbb{E}[(\ln v+\Delta c_{t+1})(1+\theta\mathbb{D}_{t+1})]-\theta\ln\delta\mathbb{E}[\mathbb{D}_{t+1}] \\
&\quad=g+\ln v+\theta(g+\ln v-\ln\delta)\Phi(z_v)-\theta\tilde{\sigma}\Phi'(z_v) \\
&\qquad-\theta\left(\sigma\omega+\ln m\right)w_v
\end{aligned} \tag{A.14}$$

Substituting (A.8) and (A.14) into (A.11) and substituting out m yields the nonlinear equation

$$1=\frac{g+\ln v+\theta\left(g+\ln\left(\frac{v}{\delta}\right)\right)\Phi(z_v)-\theta\tilde{\sigma}\Phi'(z_v)-\theta\left(\sigma\omega+\frac{1}{\rho}\ln\left(\frac{v^\rho-1}{\beta}+1\right)\right)w_v}{\frac{1}{\rho}\ln\left(\frac{v^\rho-1}{\beta}+1\right)\left(1+\theta\Phi\left(z_v\right)-\theta w_v\right)}. \tag{A.15}$$

The remaining steps mimic those for the $\alpha\neq 0$ case.

A.4 Bonds. The Euler equation for a certain payoff of one unit of the consumption good implies that the 1-period bond price equals

$$\mathcal{B}\equiv\mathbb{E}[M_{t+1}]. \tag{A.16}$$

Substituting for M_{t+1} from (10) and using Lemma 1.A.i to integrate over η^c_{t+1} yields

$$\begin{aligned}
\mathcal{B}=&\mathbb{E}\left[\beta_v\mathrm{e}^{(\alpha-1)\Delta c_{t+1}}\left(1+\theta\mathbf{1}\{\eta^c_{t+1}\leq z_v-\kappa\eta^e_{t+1}\}\right)\right] \\
=&\mathbb{E}\left[\beta_v e^{(\alpha-1)(g-\omega\sigma)+\frac{1}{2}(\alpha-1)^2\tilde{\sigma}^2+(\alpha-1)\omega\sigma\eta^e_{t+1}}\left(1+\theta\Phi(z_v-(\alpha-1)\tilde{\sigma}-\kappa\eta^e_{t+1})\right)\right].
\end{aligned} \tag{A.17}$$

Integrating over η^e_{t+1} via Lemma 2.B results in

$$\mathcal{B}=\beta_v e^{(\alpha-1)(g-\omega\sigma)+\frac{1}{2}(\alpha-1)^2\tilde{\sigma}^2}\frac{1+\theta\Phi(z_v-(\alpha-1)\tilde{\sigma})-\theta\,\mathrm{sgn}(\kappa)\exp\left\{\frac{\lambda_b}{2}-\frac{b_b}{2}\right\}\Phi\left(\lambda_b\right)}{1-(\alpha-1)\omega\sigma}. \tag{A.18}$$

where $b_b = z_v - (\alpha - 1)\tilde{\sigma}$ and $\lambda_b = \text{sgn}(\kappa)z_v - \frac{1}{|\kappa|}$. The price of a τ-period bond equals $\mathbb{E}[M_{t:t+\tau}]$, where $M_{t:t+\tau} = \prod_{h=1}^{\tau} M_{t+h}$ is the τ-period pricing kernel. Because the environment is IID, $\mathbb{E}[M_{t:t+\tau}] = \prod_{h=1}^{\tau} \mathbb{E}[M_{t:t+h}] = \mathcal{B}^{\tau}$. The yield of a τ-period bond is equal to minus the log bond price, divided by the number of periods τ,

$$-\frac{1}{\tau}\ln \mathbb{E}[M_{t:t+\tau}] = -\ln \mathcal{B}.$$

The bond yield is constant across maturities; that is, the term structure is flat as a consequence of the IID assumption.

A.5 Dividend Strips. A 1-period strip is a claim to the aggregate dividend at $t+1$, so its price is $\mathbb{E}[M_{t+1}D_{t+1}]$ and its price-dividend ratio is

$$\Psi \equiv \mathbb{E}[M_{t+1}e^{\Delta d_{t+1}}]. \tag{A.19}$$

Substituting for M_{t+1} from (10) and using Lemma 1.A.i to integrate over η^c_{t+1} and η^d_{t+1} yields

$$\begin{aligned}\Psi =& \mathbb{E}\left[\beta_v \left(1 + \theta \mathbf{1}\{\eta^c_{t+1} \leq z_v - \kappa\eta^e_{t+1}\}\right) e^{(\alpha-1)\Delta c_{t+1} + \Delta d_{t+1}}\right] \\ =& \mathbb{E}\Bigg[\beta_v \exp\left\{\alpha g - (\alpha - 1 + \varphi)\omega\sigma + \frac{1}{2}\left((\alpha-1)^2 + \varphi^2 + 2(\alpha-1)\varphi\varrho\right)\tilde{\sigma}^2\right\} \\ & \times \exp\left\{(\alpha - 1 + \varphi)\omega\sigma\eta^e_{t+1}\right\}\left(1 + \theta\Phi(z_v - (\alpha - 1 + \varrho\varphi)\tilde{\sigma} - \kappa\eta^e_{t+1})\right)\Bigg].\end{aligned} \tag{A.20}$$

Integrating over η^e_{t+1} via Lemma 2.B results in

$$\begin{aligned}\Psi =& \beta_v \exp\left\{\alpha g - (\alpha - 1 + \varphi)\omega\sigma + \frac{1}{2}\left((\alpha-1)^2 + \varphi^2 + 2(\alpha-1)\varphi\varrho\right)\tilde{\sigma}^2\right\} \\ & \times \frac{1 + \theta\Phi(b_d) - \theta\,\text{sgn}(\kappa)\exp\left\{\frac{\lambda_d^2}{2} - \frac{b_d^2}{2}\right\}\Phi(\lambda_d)}{1 - (\alpha - 1 + \varphi)\omega\sigma},\end{aligned} \tag{A.21}$$

where $b_d = z_v - (\alpha - 1 + \varrho\varphi)\tilde{\sigma}$ and $\lambda_d = \text{sgn}(\kappa)\left[\varphi(1-\varrho)\tilde{\sigma} + z_v\right] - \frac{1}{|\kappa|}$. A 1-period strip has a price of $\Psi \times D_t$ at time t and a payoff of D_{t+1} at time $t+1$, so that its return is

$$R^{(\text{1-period div strip})}_{t+1} = \frac{e^{\Delta d_{t+1}}}{\Psi}.$$

A τ-period dividend strip is a claim to the aggregate dividend at $t+\tau$, and it has a price-dividend ratio of $\Psi = \frac{1}{D_t}\mathbb{E}[M_{t:t+\tau}D_{t+\tau}]$. Due to the IID environment, $\Psi = \prod_{h=1}^{\tau}\mathbb{E}[M_{t+h}e^{\Delta d_{t+h}}] = \mathbb{E}[M_{t+1}e^{\Delta d_{t+1}}]^{\tau} = \Psi^{\tau}$. A τ-period strip has a price of $\Psi^{\tau} \times D_t$ at time t and a price of $\Psi^{\tau-1} \times D_{t+1}$ at time $t+1$, so that its holding period return is

$$R^{(\tau\text{-period div strip})}_{t+1} = \frac{e^{\Delta d_{t+1}}}{\Psi}.$$

Because the holding period return is independent of the strip's maturity, the term structure of the equity premium is flat.

A.6 Equity. Denote the price-dividend ratio of the market by $\mathcal{D}$, which is constant due to the IID setting. The cum-dividend return can be written as $R_{t+1} = e^{\Delta d_{t+1}} \frac{\mathcal{D}+1}{\mathcal{D}}$. The Euler equation for the market is given by

$$\mathcal{D} = \mathbb{E}_t \left[M_{t+1} e^{\Delta d_{t+1}} (1 + \mathcal{D}) \right]. \tag{A.22}$$

Using the definition of the price-dividend ratio of a 1-period dividend strip in (A.19), the price-dividend ratio can be expressed as

$$\mathcal{D} = \frac{\Psi}{1 - \Psi}, \tag{A.23}$$

so that the cum-dividend return equals $R_{t+1} = e^{\Delta d_{t+1}} \frac{\frac{\Psi}{1-\Psi}+1}{\frac{\Psi}{1-\Psi}} = \frac{e^{\Delta d_{t+1}}}{\Psi}$. Note that the cum-dividend return of the market equals the holding period return of dividend strips.

A.7 Forwards. The price of a 1-period forward on the dividend claim is $F_t \equiv \mathbb{E}^*[S_{t+1}]$, so the forward-to-spot ratio F_t/S_t is

$$\mathcal{F} \equiv \mathbb{E}^* \left[\frac{S_{t+1}}{S_t} \right] = \mathbb{E}^*[e^{\Delta d_{t+1}}] = \frac{1}{\mathcal{B}} \mathbb{E}[M_{t+1} e^{\Delta d_{t+1}}] = \frac{\Psi}{\mathcal{B}} \tag{A.24}$$

A.8 Options. The price of a 1-period European put option with a strike price of K equals

$$P(K) = \mathbb{E}[M_{t+1} \max\{0, K - S_{t+1}\}],$$

and it can be computed as $P(K) = \mathcal{P}(X) \times S_t$, where

$$\mathcal{P}(X) \equiv \mathbb{E} \left[M_{t+1} \max\{0, X - e^{\Delta d_{t+1}}\} \right]$$

is the put-to-spot ratio and $X \equiv K/S_t$ denotes moneyness. In what follows, I omit time-subscripts on the shocks η^c, η^d, η^e for more compact notation. Define the indicator

$$\mathbb{O}(X) \equiv \mathbf{1} \left\{ \eta^d \leq z_d - \kappa \eta^e \right\},$$

where $z_d \equiv \frac{\ln(X)-g}{\varphi\tilde{\sigma}} + \kappa$. I suppress z_d's dependence on X for notational convenience. The SDF and option payoff can be written as

$$\begin{aligned} M_{t+1} =& \beta_v e^{(\alpha-1)(g-\sigma\omega)} e^{(\alpha-1)\tilde{\sigma}\eta^c} e^{(\alpha-1)\sigma\omega\eta^e} (1 + \theta\mathbb{D}) \\ \max\{0, X - e^{\Delta d_{t+1}}\} =& \left(X - e^{g - \varphi\sigma\omega} e^{\varphi\tilde{\sigma}\eta^d} e^{\varphi\sigma\omega\eta^e} \right) \mathbb{O}(X) \end{aligned}$$

Thus,

$$\mathcal{P}(X) = \beta_v e^{(\alpha-1)(g-\sigma\omega) + \frac{1}{2}(\alpha-1)^2\tilde{\sigma}^2} \left[X \left\{ I_1(X) + \theta I_3(X) \right\} - e^{g-\varphi\sigma\omega} \left\{ I_2(X) + \theta I_4(X) \right\} \right]$$

where

$$
\begin{aligned}
I_1(X) \equiv & e^{-\frac{1}{2}(\alpha-1)^2\tilde{\sigma}^2}\mathbb{E}\left[e^{(\alpha-1)\sigma^*\eta^c}e^{(\alpha-1)\sigma\omega\eta^e}\mathbb{O}(X)\right] \\
I_2(X) \equiv & e^{-\frac{1}{2}(\alpha-1)^2\tilde{\sigma}^2}\mathbb{E}\left[e^{(\alpha-1)\sigma^*\eta^c+\varphi\sigma^*\eta^d}e^{(\alpha-1+\varphi)\sigma\omega\eta^e}\mathbb{O}(X)\right] \\
I_3(X) \equiv & e^{-\frac{1}{2}(\alpha-1)^2\tilde{\sigma}^2}\mathbb{E}\left[e^{(\alpha-1)\sigma^*\eta^c}e^{(\alpha-1)\sigma\omega\eta^e}\mathbb{O}(X)\mathbb{D}\right] \\
I_4(X) \equiv & e^{-\frac{1}{2}(\alpha-1)^2\tilde{\sigma}^2}\mathbb{E}\left[e^{(\alpha-1)\tilde{\sigma}\eta^c+\varphi\tilde{\sigma}\eta^d}e^{(\alpha-1+\varphi)\sigma\omega\eta^e}\mathbb{O}(X)\mathbb{D}\right]
\end{aligned}
$$

Note that the term $e^{\frac{1}{2}(\alpha-1)^2\tilde{\sigma}^2}$ was included in $\mathcal{P}(X)$ and its inverse in each of the four I-terms, which leads to a simplification of these terms below. Using first Lemma 1.A and then Lemma 2.B,

$$
\begin{aligned}
I_1(X) =& e^{-\frac{1}{2}(\alpha-1)^2\tilde{\sigma}^2}\mathbb{E}\left[e^{(\alpha-1)\tilde{\sigma}\eta^c}e^{(\alpha-1)\sigma\omega\eta^e}\mathbb{O}(X)\right] \\
=& \mathbb{E}\left[e^{a_1\eta^e}\Phi\left(b_1-\kappa\eta^e\right)\right] \\
=& \frac{1}{1-a_1}\left(\Phi(b_1)-\text{sgn}(\kappa)\exp\left\{\frac{\lambda_1^2}{2}-\frac{b_1^2}{2}\right\}\Phi\left(\lambda_1\right)\right)
\end{aligned}
\tag{A.25}
$$

where $a_1 \equiv (\alpha-1)\sigma\omega$, $b_1 \equiv z_d - \varrho(\alpha-1)\tilde{\sigma}$, and $\lambda_1 \equiv \frac{a_1-1+b_1\kappa}{|\kappa|}$. Using again first Lemma 1.A and then Lemma 2.B,

$$
\begin{aligned}
I_2(X) =& e^{-\frac{1}{2}(\alpha-1)^2\tilde{\sigma}^2}\mathbb{E}\left[e^{(\alpha-1)\tilde{\sigma}\eta^c+\varphi\tilde{\sigma}\eta^d}e^{(\alpha-1+\varphi)\sigma\omega\eta^e}\mathbb{O}(X)\right] \\
=& e^{\frac{1}{2}\left[\varphi^2+2\varrho(\alpha-1)\varphi\right]\tilde{\sigma}^2}\mathbb{E}\left[e^{a_2\eta^e}\Phi\left(b_2-\kappa\eta^e\right)\right] \\
=& \frac{\exp\{\frac{1}{2}\left[\varphi^2+2\varrho(\alpha-1)\varphi\right]\tilde{\sigma}^2\}}{1-a_2}\left(\Phi(b_2)-\text{sgn}(\kappa)\exp\left\{\frac{\lambda_2^2}{2}-\frac{b_2^2}{2}\right\}\Phi\left(\lambda_2\right)\right)
\end{aligned}
\tag{A.26}
$$

where $a_2 \equiv (\alpha-1+\varphi)\sigma\omega$, $b_2 \equiv z_d - \varphi\tilde{\sigma} - \varrho(\alpha-1)\tilde{\sigma}$, and $\lambda_2 \equiv \frac{a_2-1+b_2\kappa}{|\kappa|}$. Using first Lemma 1.A and then Lemma 2.E,

$$
\begin{aligned}
I_3(X) =& e^{-\frac{1}{2}(\alpha-1)^2\tilde{\sigma}^2}\mathbb{E}\left[e^{(\alpha-1)\tilde{\sigma}\eta^c}e^{(\alpha-1)\sigma\omega\eta^e}\mathbb{O}(X)\mathbb{D}\right] \\
=& \mathbb{E}\left[e^{a_3\eta^e}\Phi\left(b_3-\kappa\eta^e, d_3-\kappa\eta^e\right)\right] \\
=& \Omega(a_3, b_3, -\kappa, d_3, -\kappa)
\end{aligned}
\tag{A.27}
$$

where $a_3 \equiv (\alpha-1)\sigma\omega$, $b_3 \equiv z_d - \varrho(\alpha-1)\tilde{\sigma}$, and $d_3 \equiv \frac{x_v-g}{\tilde{\sigma}} + \kappa - (\alpha-1)\tilde{\sigma}$ and where the function Ω is defined in Lemma 2.E. Using again first Lemma 1.A and then Lemma 2.E,

$$
\begin{aligned}
I_4(X) =& e^{-\frac{1}{2}(\alpha-1)^2\tilde{\sigma}^2}\mathbb{E}\left[e^{(\alpha-1)\tilde{\sigma}\eta^c+\varphi\tilde{\sigma}\eta^d}e^{(\alpha-1+\varphi)\sigma\omega\eta^e}\mathbb{O}(X)\mathbb{D}\right] \\
=& e^{\frac{1}{2}\left[\varphi^2+2\varrho(\alpha-1)\varphi\right]\tilde{\sigma}^2}\mathbb{E}\left[e^{a_4\eta^e}\Phi\left(b_4-\kappa\eta^e, d_4-\kappa\eta^e\right)\right] \\
=& e^{\frac{1}{2}\left[\varphi^2+2\varrho(\alpha-1)\varphi\right]\tilde{\sigma}^2}\Omega(a_4, b_4, -\kappa, d_4, -\kappa)
\end{aligned}
\tag{A.28}
$$

where $a_4 \equiv (\alpha - 1 + \varphi)\sigma\omega$, $b_4 \equiv z_d - \varrho(\alpha-1)\tilde{\sigma} - \varphi\tilde{\sigma}$, and $d_4 \equiv z_v - \varrho\varphi\tilde{\sigma} - (\alpha-1)\tilde{\sigma}$ and where the function Ω is defined in Lemma 2.E. The call-to-spot ratio $\mathcal{C}(X)$ follows from the normalized put-call parity,

$$\mathcal{C}(X) = \mathcal{P}(X) + \mathcal{B}(\mathcal{F} - X), \tag{A.29}$$

making it unnecessary to price calls separately.

A.9 VIX. The squared VIX index is defined as

$$VIX_t^2 \equiv 2R_{t+1}^f \left(\int_0^{F_t} \frac{P_t(K)}{K^2} dK + \int_{F_t}^{\infty} \frac{C_t(K)}{K^2} dK \right). \tag{A.30}$$

The official VIX index, computed by the Chicago Board Options Exchange (CBOE), is based on a discretized version of this equation that replaces the integrals by sums. In the model, the VIX can be computed without finding option prices by first re-writing it as the risk-neutral entropy of returns. To do so, let S_t denote the ex-dividend price of the market. Starting from the identity[17]

$$\int_0^s \frac{\max\{X - S_{t+1}, 0\}}{X^2} dX + \int_s^{\infty} \frac{\max\{S_{t+1} - X, 0\}}{X^2} dX = \left(\frac{S_{t+1}}{s} - 1 \right) - \ln\left(\frac{S_{t+1}}{s} \right), \tag{A.31}$$

set s equal to the forward price $\mathbb{E}_t^*[S_{t+1}]$ and take the risk-neutral expectation, which yields

$$2R_{t+1}^f \left(\int_0^{F_t} \frac{P_t(X)}{X^2} dX + \int_{F_t}^{\infty} \frac{C_t(X)}{X^2} dX \right) = \ln \mathbb{E}_t^*[S_{t+1}] - \mathbb{E}_t^* \left[\ln S_{t+1} \right], \tag{A.32}$$

where I have used that $P_t(X) = \frac{\mathbb{E}^*[\max\{0, X - S_{t+1}\}]}{R_{t+1}^f}$ and $C_t(X) = \frac{\mathbb{E}^*[\max\{0, S_{t+1} - X\}]}{R_{t+1}^f}$. Finally, adding and subtracting $\ln S_t$ on the RHS and multiplying by 2 yields

$$VIX_t^2 = 2\left(\ln \mathbb{E}_t^* \left[\frac{S_{t+1}}{S_t} \right] - \mathbb{E}_t^* \left[\ln\left(\frac{S_{t+1}}{S_t} \right) \right] \right), \tag{A.33}$$

that is, the squared VIX index equals 2 times the risk-neutral conditional entropy of the ex-dividend market return. In the IID setting considered here, there is no state dependence and the ex-dividend return equals $\frac{S_{t+1}}{S_t} = e^{\Delta d_{t+1}}$, so that the squared VIX is

$$VIX_t^2 = 2\left(\ln \mathbb{E}^* \left[e^{\Delta d_{t+1}} \right] - \mathbb{E}^* \left[\Delta d_{t+1} \right] \right). \tag{A.34}$$

The first term in (A.34) is given by $\mathbb{E}^* \left[e^{\Delta d_{t+1}} \right] = \mathcal{F}$. To find the second term, note that $\mathbb{E}^* \left[\Delta d_{t+1} \right] = \frac{1}{\mathcal{B}} \mathbb{E} \left[M_{t+1} \Delta d_{t+1} \right]$ and define $\Gamma \equiv \mathbb{E} \left[M_{t+1} \Delta d_{t+1} \right]$. Substituting for M_{t+1} from (10) and using Lemmas 1.A.1 and 1.C to integrate over η_{t+1}^c yields

[17] The easiest way to see that (A.31) is an identity is to evaluate it separately for $S_{t+1} \leq s$ and $S_{t+\tau} \geq s$.

$$
\begin{aligned}
\Gamma =&\mathbb{E}\left[\beta_v e^{(\alpha-1)\Delta c_{t+1}}\left(1+\theta\mathbf{1}\{\eta^c_{t+1}\le z_v-\kappa\eta^e_{t+1}\}\right)\Delta d_{t+1}\right]\\
=&\beta_v e^{(\alpha-1)(g-\sigma\omega)}\mathbb{E}\Bigg[e^{(\alpha-1)(\tilde\sigma\eta^c_{t+1}+\sigma\omega\eta^e_{t+1})}\left(1+\theta\mathbf{1}\{\eta^c_{t+1}\le z_v-\kappa\eta^e_{t+1}\}\right)\\
&\times\left(g-\varphi\sigma\omega+\varphi\tilde\sigma\eta^d_{t+1}+\varphi\sigma\omega\eta^e_{t+1}\right)\Bigg]\\
=&\beta_v e^{(\alpha-1)(g-\sigma\omega)}\Bigg\{\mathbb{E}\left[e^{(\alpha-1)(\tilde\sigma\eta^c_{t+1}+\sigma\omega\eta^e_{t+1})}\left(g-\varphi\sigma\omega+\varphi\tilde\sigma\eta^d_{t+1}+\varphi\sigma\omega\eta^e_{t+1}\right)\right]\\
&+\theta\mathbb{E}\left[e^{(\alpha-1)(\tilde\sigma\eta^c_{t+1}+\sigma\omega\eta^e_{t+1})}\mathbf{1}\{\eta^c_{t+1}\le z_v-\kappa\eta^e_{t+1}\}\left(g-\varphi\sigma\omega+\varphi\sigma\omega\eta^e_{t+1}\right)\right]\Bigg\}\\
&+\theta\varphi\tilde\sigma\mathbb{E}\left[e^{(\alpha-1)(\tilde\sigma\eta^c_{t+1}+\sigma\omega\eta^e_{t+1})}\mathbf{1}\{\eta^c_{t+1}\le z_v-\kappa\eta^e_{t+1}\}\eta^d_{t+1}\right]\Bigg\} \qquad \text{(A.35)}\\
=&\beta_v e^{(\alpha-1)(g-\sigma\omega)+\frac{1}{2}(\alpha-1)^2\tilde\sigma^2}\Bigg\{\mathbb{E}\left[e^{(\alpha-1)\sigma\omega\eta^e_{t+1}}\left(g-\varphi\sigma\omega+\varrho(\alpha-1)\varphi\tilde\sigma^2+\varphi\sigma\omega\eta^e_{t+1}\right)\right]\\
&+\theta\mathbb{E}\left[e^{(\alpha-1)\sigma\omega\eta^e_{t+1}}\Phi(b-\kappa\eta^e_{t+1})\left(g-\varphi\sigma\omega+\varphi\sigma\omega\eta^e_{t+1}\right)\right]\\
&+\theta\varphi\tilde\sigma\varrho\mathbb{E}\left[e^{(\alpha-1)\sigma\omega\eta^e_{t+1}}\left((\alpha-1)\tilde\sigma\Phi(b-\kappa\eta^e_{t+1})-\frac{1}{\sqrt{2\pi}}\mathrm{e}^{-\frac{(z_v-\kappa\eta^e_{t+1}-(\alpha-1)\tilde\sigma)^2}{2}}\right)\right]\Bigg\}\\
=&\beta_v e^{(\alpha-1)(g-\sigma\omega)+\frac{1}{2}(\alpha-1)^2\tilde\sigma^2}\Bigg\{\mathbb{E}\left[e^{a\eta^e_{t+1}}\left(1+\theta\Phi(b-\kappa\eta^e_{t+1})\right)\left(d+\varphi\sigma\omega\eta^e_{t+1}\right)\right]\\
&-\frac{\theta\varphi\tilde\sigma\varrho}{\sqrt{2\pi}}\mathbb{E}\left[e^{a\eta^e_{t+1}-\frac{1}{2}(b-\kappa\eta^e_{t+1})^2}\right]\Bigg\}
\end{aligned}
$$

where the last step introduced the simplifying notation

- $a_x \equiv (\alpha-1)\omega\sigma$
- $b_x \equiv z_v-(\alpha-1)\tilde\sigma$
- $d_x \equiv g-\varphi\omega\sigma+(\alpha-1)\varrho\varphi\tilde\sigma^2$

Additionally defining $\lambda_x \equiv \operatorname{sgn}(\kappa)z_v-\frac{1}{|\kappa|}$ and using Lemmas 2.A, 2.B, and 2.D to Integrate over η^e yields

$$
\begin{aligned}
\Gamma =&\beta_v e^{(\alpha-1)(g-\sigma\omega)+\frac{1}{2}(\alpha-1)^2\tilde\sigma^2}\\
&\times\Bigg\{\frac{d_x+\varphi\sigma\omega/(1-a_x)}{1-a_x}\left(1+\theta\Phi(b_x)-\theta\operatorname{sgn}(\kappa)\mathrm{e}^{\frac{\lambda_x^2}{2}-\frac{b_x^2}{2}}\Phi(\lambda_x)\right)\\
&-\frac{\theta\varphi\tilde\sigma}{1-a_x}\left(\Phi'(b_x)+\lambda\mathrm{e}^{\frac{\lambda_x^2}{2}-\frac{b_x^2}{2}}\Phi(\lambda_x)\right)-\frac{\theta\varphi\tilde\sigma\varrho}{|\kappa|}\mathrm{e}^{\frac{1}{2}z_v^2-\frac{z_v}{\kappa}+\frac{1}{2\kappa^2}}\Phi\left(\lambda_x\right)\Bigg\}
\end{aligned} \qquad \text{(A.36)}
$$

Substituting back into (A.34) shows that the squared VIX is given by

$$
VIX^2 = 2\left(\ln\mathcal{F}-\Gamma/\mathcal{B}\right). \qquad \text{(A.37)}
$$

A.10 SVIX. Following Martin (2017), the squared simple VIX (SVIX) is defined as

$$SVIX_t^2 \equiv \frac{2}{R_{t+1}^f}\left[\int_0^{F_t}\frac{P_t(K)}{S_t^2}dK + \int_{F_t}^{\infty}\frac{C_t(K)}{S_t^2}dK\right]. \tag{A.38}$$

Whereas VIX^2 in (A.30) weights option prices by the inverse of their squared strike price, $SVIX^2$ represents an equally weighted average. Martin (2017) shows that, if the underlying does not pay a dividend, $SVIX^2$ measures $\frac{1}{(R_{t+1}^f)^2}$ times the risk-neutral variance of the underlying's simple return (see pp. 379–82 of Martin's study). The steps of the proof also remain valid, however, when the underlying does pay a dividend. In that case, $SVIX^2$ measures $\frac{1}{(R_{t+1}^f)^2}$ times the risk-neutral variance of the underlying's simple ex-dividend return. Equation 13 in Martin (2017), expressed in the notation used here and simplified to the IID setting is

$$SVIX^2 = \mathcal{B}^2\mathbb{E}^*[e^{2\Delta d_{t+1}}] - \Psi^2 \tag{A.39}$$

To find the remaining term, note that $\mathbb{E}^*[e^{2\Delta d_{t+1}}] = \frac{1}{\mathcal{B}}\mathbb{E}[M_{t+1}e^{2\Delta d_{t+1}}]$ and define $\Lambda \equiv \mathbb{E}[M_{t+1}e^{2\Delta d_{t+1}}]$. Substituting for M_{t+1} from (10) and using Lemma 1.A.1 to integrate over η^c yields

$$\begin{aligned}\Lambda =&\mathbb{E}\left[\beta_v e^{(\alpha-1)\Delta c_{t+1}+2\Delta d_{t+1}}\left(1+\theta\mathbf{1}\{\eta_{t+1}^c \le z_v - \kappa\eta_{t+1}^e\}\right)\right]\\ =&\mathbb{E}\left[\beta_v \exp\left((\alpha+1)g-(\alpha-1+2\varphi)\omega\sigma + 2\tilde{\sigma}^2\left(\frac{1}{4}(\alpha-1)^2+\varphi^2+(\alpha-1)\varrho\varphi\right)\right)\right.\\ &\left.\times e^{(\alpha-1+2\varphi)\omega\sigma\eta_{t+1}^e}\left(1+\theta\Phi(z_v-(\alpha-1+2\varrho\varphi)\tilde{\sigma}-\kappa\eta_{t+1}^e)\right)\right]\end{aligned} \tag{A.40}$$

Integrating over η_{t+1}^e via Lemma 2.B results in

$$\begin{aligned}\Lambda =&\beta_v \exp\left((\alpha+1)g-(\alpha-1+2\varphi)\omega\sigma + 2\tilde{\sigma}^2\left(\frac{1}{4}(\alpha-1)^2+\varphi^2+(\alpha-1)\varrho\varphi\right)\right)\\ &\times\frac{1+\theta\Phi(b_s)-\operatorname{sgn}(\kappa)\theta\exp\left(\frac{\lambda_s^2}{2}-\frac{b_s^2}{2}\right)\Phi(\lambda_s)}{1-(\alpha-1+2\varphi)\omega\sigma}\end{aligned} \tag{A.41}$$

where $b_s = z_v - (\alpha-1+2\varrho\varphi)\tilde{\sigma}$ and $\lambda_s = \operatorname{sgn}(\kappa)z_v - \frac{1}{|\kappa|} + 2\tilde{\sigma}\varphi(1-\varrho)$. Substituting back into (A.39) shows that the squared SVIX is given by

$$SVIX^2 = \mathcal{B}\Lambda - \Psi^2. \tag{A.42}$$

References

Abel, A. B. (1999): "Risk premia and term premia in general equilibrium," *Journal of Monetary Economics*, 43, 3–33.

Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi (2017): "Microeconomic Origins of Macroeconomic Tail Risks," *American Economic Review*, 107, 54–108.

Ait-Sahalia, Y. and A. W. Lo (2000): "Nonparametric risk management and implied risk aversion," *Journal of Econometrics*, 94, 9–51.

Allais, M. (1979): "The foundations of a positive theory of choice involving risk and a criticism of the postulates and axioms of the American school," in *Expected Utility Hypothesis and the Paradox*, ed. by M. Allais and O. Hagon, Dordrecht, Holland: D. Reidel Publishing Co.

Andries, M. and V. Haddad (2019): "Information Aversion," Working Paper.

Ang, A. and J. Chen (2002): "Asymmetric correlations of equity portfolios," *Journal of Financial Economics*, 63, 443–494.

Ang, A., J. Chen, and Y. Xing (2006): "Downside Risk," *Review of Financial Studies*, 19, 11911239.

Backus, D., M. Chernov, and I. Martin (2011): "Disasters Implied by Equity Index Options," *Journal of Finance*, 66, 1969–2012.

Backus, D. K., B. R. Routledge, and S. E. Zin (2010): "The cyclical component of US asset returns," Working Paper.

Bansal, R., R. F. Dittmar, and C. T. Lundblad (2005): "Consumption, Dividends, and the Cross Section of Equity Returns," *Journal of Finance*, 60, 1639–1672.

Bansal, R., D. Kiku, and A. Yaron (2012): "An Empirical Evaluation of the Long-Run Risks Model for Asset Prices," *Critical Finance Review*, 1, 183–221.

Bansal, R. and A. Yaron (2004): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 57, 1481–1509.

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

——— (2009): "Rare Disasters, Asset Prices, and Welfare Costs," *American Economic Review*, 99, 243–264.

Bates, D. S. (2008): "The market for crash risk," *Journal of Economic Dynamics and Control*, 32, 2291–2321.

Beason, T. and D. Schreindorfer (2019): "On Sources of Risk Premia in Representative Agent Models," Working Paper.

Beeler, J. and J. Y. Campbell (2012): "The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment," *Critical Finance Review*, 1, 141–182.

Bekaert, G. and E. Engstrom (2017): "Asset Return Dynamics under Habits and Bad Environment-Good Environment Fundamentals," *Journal of Political Economy*, 60, 713–760.

Berger, D., I. Dew-Becker, and S. Giglio (2019): "Uncertainty shocks as second-moment news shocks," *Review of Economic Studies*, forthcoming.

Bhamra, H. S., L.-A. Kuehn, and I. A. Strebulaev (2010): "The Levered Equity Risk Premium and Credit Spreads: A Unified Framework," *Review of Financial Studies*, 23, 645703.

Black, F. and M. Scholes (1973): "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 637–654.

Bollerslev, T., G. Tauchen, and H. Zhou (2009): "Expected Stock Returns and Variance Risk Premia," *Review of Financial Studies*, 4464–4492.

Bollerslev, T. and V. Todorov (2011): "Tails, fears, and risk premia," *Journal of Finance*, 66, 2165–2211.

Bonomo, M., R. Garcia, N. Meddahi, and R. Tedongap (2011): "Generalized Disappointment Aversion, Long-run Volatility Risk, and Asset Prices," *Review of Financial Studies*, 24, 82–122.

Breeden, D. and R. H. Litzenberger (1978): "State Contingent Prices Implicit in option Prices," *Journal of Business*, 51, 3–24.

Broadie, M., M. Chernov, and M. Johannes (2009): "Understanding Index Option Returns," *Review of Financial Studies*, 22, 4493–4529.

Campanale, C., R. Castro, and G. L. Clementi (2010): "Asset Pricing in a Production Economy with Chew-Dekel Preferences," *Review of Economic Dynamics*, 13, 379–402.

Campbell, J. Y. and J. H. Cochrane (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205–251.

Chen, L. (2009): "On the reversal of return and dividend growth predictability: A tale of two periods," *Journal of Financial Economics*, 92, 128–151.

Coval, J. A. and T. Shumway (2001): "Expected Option Returns," *Journal of Finance*, 54, 983–1009.

Dahlquist, M., A. Farago, and R. Tedongap (2016): "Asymmetriies and Portfolio Choice," *Review of Financial Studies*, 30, 667–702.

Delikouras, S. (2017): "Wheres the Kink? Disappointment Events in Consumption Growth and Equilibrium Asset Prices," *Review of Financial Studies*, 30, 28512889.

Delikouras, S. and A. Kostakis (2019): "A Single-Factor Consumption-Based Asset Pricing Model," *Journal of Financial and Quantitative Analysis*, 54, 789–827.

Dew-Becker, I., S. Giglio, A. Le, and M. Rodriquez (2017): "The price of variance risk," *Journal of Financial Economics*, 123, 225–250.

Drechsler, I. (2013): "Uncertainty, Time-Varying Fear, and Asset Prices," *Journal of Finance*, 68.

Drechsler, I. and A. Yaron (2010): "Whats Vol Got to Do with It," *Review of Financial Studies*, 24.

Duffee, G. R. (2005): "Time Variation in the Covariance between Stock Returns and Consumption Growth," *Journal of Finance*, 60, 1673–1712.

Epstein, L. G., E. Farhi, and T. Strzalecki (2014): "How much would you pay to resolve long-run risk," *American Economic Review*, 104, 2680–2697.

Epstein, L. G. and S. E. Zin (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937–969.

——— (1991): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis," *Journal of Political Economy*, 99, 263–286.

——— (2001): "The Independence Axiom and Asset Returns," *Journal of Empirical Finance*, 8, 537–572.

Farago, A. and R. Tedongap (2018): "Downside risks and the cross-section of asset returns," *Journal of Financial Economics*, 129, 69–86.

Gabaix, X. (2012): "Variable Rare Disasters: An Exactly Solved Framework for ten Puzzles in Macro-finance," *Quarterly Journal of Economics*, 127, 645–700.

Gul, F. (1991): "A Theory of Disappointment Aversion," *Econometrica*, 59, 667–686.

Kroencke, T. A. (2017): "Asset Pricing without Garbage," *Journal of Finance*, 70, 47–98.

Lettau, M., M. Maggioro, and M. Weber (2014): "Conditional risk premia in currency markets and other asset classes," *Journal of Financial Economics*, 114, 197–225.

Liu, H. and J. Miao (2015): "Growth uncertainty,g eneralized disappointment aversion and production-based asset pricing," *Journal of Monetary Economics*, 69, 70–89.

Longin, F. and B. Solnik (2001): "Extreme Correlation of International Equity Markets," *Journal of Finance*, 56, 649–676.

Lucas, R. E. J. (1987): *Models of Business Cycles*, Basil Blackwell, New York.

Martin, I. (2017): "What is the expected return on the market," *Quarterly Journal of Economics*, 132, 367–433.

Mehra, R. and E. C. Prescott (1985): "The equity premium: A puzzle," *Journal of Monetary Economics*, 15, 145–161.

Rietz, T. A. (1988): "The equity risk premium: a solution," *Journal of Monetary Economics*, 22, 117–131.

Rosenberg, J. V. and R. F. Engle (2002): "Empirical pricing kernels," *Journal of Financial Economics*, 64, 341–372.

Routledge, B. R. and S. E. Zin (2010): "Generalized Disappointment Aversion and Asset Prices," *Journal of Finance*, 65, 1303–1332.

Schreindorfer, D. (2014): "Tails, Fears, and Equilibrium Option Prices," Working Paper.

Wachter, J. (2013): "Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?" *Journal of Finance*, 68, 987–1035.

Weil, P. (1989): "The equity premium puzzle and the risk-free rate puzzle," *Journal of Monetary Economics*, 24, 401–421.

Working, H. (1960): "Note on the Correlation of First Differences of Aggregates in a Random Chain," *Econometrica*, 28, 916–918.

Xu, N. (2018): "Procyclicality of the comovement between dividend growth and consumption growth," Working Paper.