

Generalized Disappointment Aversion and the Variance Term Structure

Mykola Babiak 
Lancaster University Management School
m.babiak@lancaster.ac.uk

Abstract

Contrary to leading asset pricing theories, recent empirical evidence indicates that financial markets compensate only short-term equity variance risk. An equilibrium model with generalized disappointment aversion risk preferences and rare events reconciles salient features of the variance term structure. In addition, a calibration explains the variance and skew risk premiums in equity returns and the implied volatility skew of index options while capturing standard moments of fundamentals, equity returns, and the risk-free rate. The key intuition for the results stems from substantial countercyclical risk aversion induced by endogenous variation in the probability of disappointing events in consumption growth.

I. Introduction

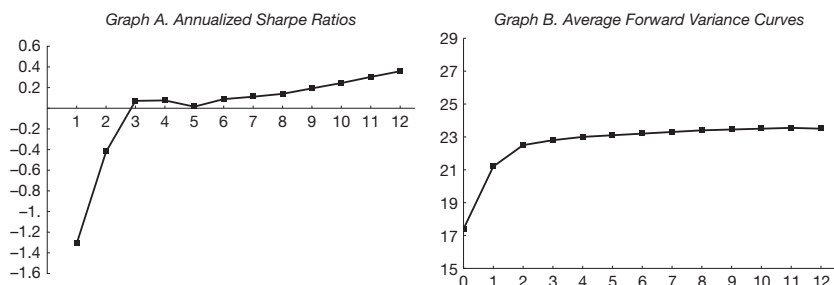
The consumption-based asset pricing literature has been recently revived by generalized models of long-run risks and rare disasters to capture many characteristics of the equity and derivatives markets. Yet leading theories fail to explain the timing of variance risk. Contrary to most successful asset pricing models, Dew-Becker, Giglio, Le, and Rodriguez (2017) show that it has been costless to hedge future variance at horizons longer than 2 months, whereas only unexpected

This article is based on the first chapter of my dissertation at CERGE-EI (2019) and was previously circulated under the title “Generalized Disappointment Aversion, Learning, and Asset Prices.” I am especially indebted to Roman Kozhan for constructive suggestions that greatly improved the article. I appreciate helpful comments from Anmol Bhandari, Daniele Bianchi, Jaroslav Borovicka, Thierry Foucault (the editor), Michael Hasler, Michael Johannes, Kenneth L. Judd, Marek Kapicka, Michal Kejak, Michal Pakoš, Bryan Routledge (the referee), David Schreindorfer, Veronika Selezneva, Ctirad Slavik, Sergey Slobodyan, Stijn Van Nieuwerburgh, Ansgar Walther, conference participants at the 2016 EEA-ESEM Meeting, the 2016 Meeting of the Society for Computational Economics, the 2016 Zurich Initiative for Computational Economics, the 2016 Annual Conference of the Swiss Society for Financial Market Research, the 2018 RES Annual Conference, the 2018 RES Symposium of Junior Researchers, the 2018 Spanish Economic Association Meeting, and seminar participants at Columbia Business School, Warwick Business School, Lancaster University Management School, Collegio Carlo Alberto, Durham University Business School, the University of Gothenburg, and the University of Groningen. The financial support from the Charles University Grant Agency (GAUK No. 151016) and the Czech Science Foundation Project No. P402/12/G097 (DYME Dynamic Models in Economics) is gratefully acknowledged.

FIGURE 1

Average Prices and Annualized Sharpe Ratios for Forward Variance Claims

Figure 1 plots annualized Sharpe ratios and average prices for forward variance claims in the U.S. data from 1996 to 2013. The prices are reported in annualized volatility terms. The data are from Dew-Becker et al. (2017).



realized variance was significantly priced.¹ The term structure of variance risk possess a challenge to models featuring time-varying expected growth and volatility (Bansal and Yaron (2004)) or disaster risk (Rietz (1988), Barro (2006)).

Illustrate the challenge in Figure 1 by showing the empirical Sharpe ratios and prices for forward variance claims, which are swap contracts that pay the owner the realized stock market variance during a particular future period.² The figure shows the term structure of forward claims on future variance up to 1 year. The average prices are upward-sloping at the short end and quickly flatten with the horizon. Sharpe ratios are significantly negative for short maturities, suggesting investors are willing to hedge short-term variance risk. Puzzling, however, is that future variance from 3 to 12 months is unpriced. Well-known asset pricing theories predict a strongly upward-sloping term structure of forward variance prices and, hence, imply the negative and significant Sharpe ratios at future horizons, counter to what we observe empirically.

I capture the observed variance term structure by introducing asymmetric preferences into a model with learning about consumption depressions.³ Disaster risk generates the upward-sloping term structure of return variance, however, I demonstrate that asymmetric preferences cancel the increasing effect in the long term. The reason is that, in bad times, forward variance becomes higher in the short term than in the long term with asymmetric preferences, which flattens the increasing

¹Dew-Becker, Giglio, and Kelly (2021) show that it is highly costly to hedge realized volatility but not forward-looking uncertainty across different markets. Berger, Dew-Becker, and Giglio (2019) provide new empirical evidence that shocks to future uncertainty have no significant effect on the economy. Also, Dew-Becker and Giglio (2020) find that investors do not view shocks to cross-sectional uncertainty as bad. Also, van Binsbergen, Brandt, and Kojien (2012) and van Binsbergen, Hueskes, Kojien, and Vrugt (2013) document a downward-sloping term structure of equity risk premia and volatility, which is at odds with leading asset pricing models.

²For instance, a payoff (realized variance) of n -month variance forward equals the sum of daily squared stock market returns in month n from today.

³The ingredients are empirically motivated. A number of studies provide micro-level evidence that investors dislike losses more than they enjoy gains (Choi, Fisman, Gale, and Kariv (2007)). Also, Hansen (2007) argues that the assumption of the investor's full information about the model structure is extreme.

pattern at longer horizons on average. The properties of forward return variance translate into empirically consistent variance forward prices. This mechanism also implies negative Sharpe ratios on short-term variance forwards and positive and increasing ratios at longer maturities.

Formally, I consider an exchange economy with generalized disappointment aversion (GDA) risk preferences (Routledge and Zin (2010)) and rare events. Consumption growth follows a hidden two-state Markov chain where a rare “depression” is calibrated to the U.S. Great Depression. The agent filters the hidden state probabilities. GDA preferences amplify the impact on the pricing kernel of disappointing beliefs corresponding to utilities below a scaled certainty equivalent. The amplification of lower-tail shocks yields strongly countercyclical risk aversion, which helps capture the variance term structure.

The economic mechanism is as follows: Following Veronesi (1999), the conditional volatility of equity return is a hump-shaped function of a posterior probability of expansion, π_t (GDA in Figure 4). The economy is in a good state for most of the periods, in which case π_t is high and close to 1. A good piece of news reinforces the investor’s beliefs that the current regime is the expansion. In this case, the risk of future disasters generates an upward-sloping term structure of forward variance. A bad piece of news decreases π_t and leads to a spike in return variance initially. Bad news could be due to a disaster and hence the investor will learn times are bad in the future ($\pi_t \approx 0$). Bad innovations could also be due to idiosyncratic consumption risk in expansion and hence the investor will update beliefs to reflect times are still good ($\pi_t \approx 1$). In both cases, return variance will decrease quickly when π_t approaches 0 or 1, implying the inversion in forward variance. Unconditionally, the investor is always willing to hedge high realized variance in the short term. In the long term, however, the inversion in bad times dominates the upward-sloping effect of disaster risk in good periods, flattening the forward variance curve. Variance claims inherit the properties of forward variance. Thus, the unconditional term structure of prices is upward-sloping at the short end and flattens out quickly in maturity. The inversion in prices yields positive Sharpe ratios on variance forwards at longer horizons on average.

Intuitively, the inversion in forward variance in response to bad news happens because high volatility is short-lived in the economy.⁴ Indeed, the conditional volatility peaks within a narrow range of beliefs and sharply diminishes outside this interval. When beliefs change, return volatility spikes but does not persist. Mechanically, sizable countercyclical risk aversion induced by GDA preferences yields strong and weak price sensitivities to belief changes in good and bad times (Veronesi (1999)). This difference in sensitivities implies return volatility should be higher following a bad piece of news in good times than a good piece of news in bad times. As a result, an asymmetric effect on the price sensitivity to news leads to a skewed shape of conditional volatility.

Next, I compare GDA preferences with nested utility functions. I show that the term structure of variance risk can be replicated with GDA preferences due

⁴This mechanism is consistent with Dew-Becker et al. (2017) showing that, during consumption disasters and financial crises, realized volatility spikes for 1 month only and then reverts quickly.

to a sufficiently countercyclical risk aversion.⁵ Interestingly, not only can nested preference specifications be rejected by the unconditional term structure, but they are also inconsistent with the conditional dynamics of the variance term structure.

I first compare GDA preferences to a disappointment aversion utility function (Gul (1991)) and Epstein–Zin preferences (Epstein and Zin (1989)). First, a disappointment-averse agent increases the pricing kernel for disappointing utilities, defined as being below the certainty equivalent. Compared to Routledge and Zin (2010), Gul's preferences increase the disappointment threshold. This generates a large number of disappointing events and a large risk aversion in two states. Thus, price sensitivities are similar in good and bad times, generating a symmetric shape of return volatility (DA in Figure 4). Second, a model with Epstein–Zin preferences generates a slightly skewed shape of return volatility (EZ in Figure 4). However, conditional volatility remains elevated for a wide range of beliefs in both models. When the investor's beliefs change, high variance persists in the long term. This generates the upward-sloping term structure of forward variance and prices.

I also look at the conditional dynamics of the term structures. I assume the investor holds a median belief (normal times). I then study the impact of one positive and three negative consumption innovations. First, at the 1-month maturity, the average Sharpe ratios in the GDA economy are pro-cyclical, meaning more (less) negative in bad (good) times, consistent with Ait-Sahalia, Karaman, and Mancini (2020). The reason is that GDA preferences generate a beliefs-dependent pricing kernel with higher marginal utility in low consumption states, increasing the hedge against high realized variance associated with low-utility states.⁶ At longer maturities, the Sharpe ratios remain close to 0 in response to small shocks, whereas they become upward-sloping and positive in response to large negative news. The reason is that small shocks are not priced due to a low disappointment threshold. In contrast, lower-tail shocks place the posterior belief within the interval of the highest return variance and, hence, the variance tends to be lower in later periods. The variance claims are priced accordingly, making the short-term variance forward more expensive. The inversion in prices generates positive Sharpe ratios at longer horizons.

Second, in the disappointment aversion model, the conditional variance forward prices remain strongly upward-sloping, implying negative Sharpe ratios across all economic conditions. The reason is that high variance is persistent due to the shape of the conditional return variance and, therefore, variance risk concentrates in the long term. Third, in the Epstein–Zin economy, prices of variance forwards remain markedly increasing in the horizon for most economic conditions and become mildly decreasing only when consumption growth is extremely low. The mild inversion is too weak to dampen the upward-sloping effect at other times. Thus, the conditional Sharpe ratios remain strongly negative.

⁵The countercyclical risk aversion can rationalize the equity premium puzzle (Melino and Yang (2003)). I show that, in my setting, a sufficiently countercyclical risk aversion induced by generalized disappointment aversion can further explain the variance term structure.

⁶Routledge and Zin (2010) and Bonomo, Garcia, Meddahi, and Tédongap (2011) provide a similar analysis of GDA stochastic discount factor with alternative consumption processes. Also, the beliefs-dependent effective risk aversion of my paper echoes the mechanism of Berrada, Detemple, and Rindisbacher (2018) with learning and a beliefs-dependent utility function.

Finally, the GDA model shows superior performance when confronted with other asset pricing facts. It captures salient features of the equity variance and skew risk premiums and a volatility skew implied by index option prices.⁷ In contrast, other frameworks generate too small variance and skew risk premiums and flat implied volatility curves. In a comparative analysis, I show that my results are robust to different calibrations of key parameter values. Following Pohl, Schmedders, and Wilms (2018) and Lorenz, Schmedders, and Schumacher (2020), I check that global projection methods provide highly accurate numerical solutions.

This article is related to several strands of the literature. First, it contributes to the growing literature on the term structures of equity and variance claims (van Binsbergen et al. (2012), (2013), Dew-Becker et al. (2017)). A number of studies (Croce, Lettau, and Ludvigson (2014), Belo, Collin-Dufresne, and Goldstein (2015), Favilukis and Lin (2015), Hasler and Márfe (2016), Márfe (2017), Ai, Croce, Diercks, and Li (2018), and Hasler, Khapko, and Máfe (2019)) explain the downward-sloping term structure of equity risk premia and return volatility.⁸ I complement these articles by explaining the variance term structure.

Second, this study builds on the literature exploring asset pricing properties of GDA preferences. These preferences have been used to explain stock market returns (Bonomo et al. (2011), Bonomo, Garcia, Meddahi, and Tédongap (2015), Liu and Miao (2014), and Schreindorfer (2020)), sovereign spreads (Augustin and Tédongap (2016)), portfolios (Dahlquist, Farago, and Tédongap (2016)), the cross section of stock returns (Delikouras (2017), Farago and Tédongap (2018), and Delikouras and Kostakis (2019)), and the term structure of interest rates (Augustin and Tédongap (2021)). I employ GDA preferences to explain the variance forward prices and returns. This article is, to my knowledge, the first to reconcile the variance term structure. It does so while jointly explaining equity returns, variance and skew premiums, and option prices. Also, the extant literature studies GDA preferences in long-run risk models, while this article examines a rare event model with learning.

Third, this article is related to leading asset pricing theories focusing on the variance premium and option prices. These include the extensions of equilibrium models with habit (Du (2011)), rare disasters (Liu, Pan, and Wang (2005), Benzoni, Collin-Dufresne, and Goldstein (2011), and Sco and Wachter (2019)), and long-run risks (Eraker and Shaliastovich (2008), Bollerslev, Tauchen, and Zhou (2009), Drechsler and Yaron (2011), Drechsler (2013), Zhou and Zhu (2014), and Shaliastovich (2015)). My article is distinct from this literature because it points out the importance of the investor's GDA for the variance term structure.

Finally, this article connects to hidden Markov switching models (David (1997), Veronesi (1999), (2000)).⁹ The recent literature extends this approach to learning about unknown volatility (Weitzman (2007)) and persistence (Cogley and Sargent (2008), Gillman, Kejak, and Pakos (2015), and Andrei, Hasler, and Jeanneret (2019)) as well as to a multidimensional-learning problem (Collin-Dufresne, Johannes,

⁷Also, see Choi, Mueller, and Vedolin (2017) and Londono and Zhou (2017) for bond and currency variance risk premiums.

⁸See van Binsbergen and Kojen (2017) for a review of the literature on term structures of equity claims.

⁹See Pastor and Veronesi (2009) for a survey of the early literature on learning in financial markets.

and Lochstoer (2016), Johannes, Lochstoer, and Mou (2016), and Babiak and Kozhan (2020), (2021)). This article contributes to the learning literature by investigating how state uncertainty is priced in the presence of GDA preferences with a particular emphasis on the pricing of the variance risk.

The remainder of the article is organized as follows: Section II describes the economy. Section III outlines the equilibrium conditions. Section IV provides asset pricing results and sensitivity analysis. Section V concludes. Supplementary Material provides supporting analysis and additional results.

II. Model

A. Generalized Disappointment Aversion Risk Preferences

The environment is an infinite-horizon, discrete-time exchange economy with a representative agent. Following Epstein and Zin (1989), the agent's utility V_t is defined by

$$(1) \quad V_t = [(1 - \beta)C_t^\rho + \beta\mathcal{R}_t^\rho]^{1/\rho},$$

in which C_t is consumption, $0 < \beta < 1$ is the subjective discount factor, $1/(1 - \rho) > 0$ is the elasticity of intertemporal substitution (EIS), and $\mathcal{R}_t = \mathcal{R}_t(V_{t+1})$ is the certainty equivalent.

The certainty equivalent captures the GDA risk of Routledge and Zin (2010). GDA preferences put more weight on “disappointing” events compared to the expected utility, similar to disappointment aversion risk preferences of Gul (1991). For Gul's model, however, an outcome is viewed as disappointing when it is below the certainty equivalent, whereas for Routledge and Zin's specification a disappointing outcome is below a constant fraction of the implicit certainty equivalent. Formally, the certainty equivalent of GDA preferences is implicitly defined by

$$(2) \quad \frac{[\mathcal{R}_t(V_{t+1})]^\alpha}{\alpha} = \mathbb{E}_t \left[\frac{V_{t+1}^\alpha}{\alpha} \right] - \theta \mathbb{E}_t \left[\mathbb{I} \left(\frac{V_{t+1}}{\mathcal{R}_t(V_{t+1})} \leq \delta \right) \left(\frac{[\delta \mathcal{R}_t(V_{t+1})]^\alpha}{\alpha} - \frac{V_{t+1}^\alpha}{\alpha} \right) \right],$$

in which $\mathbb{I}(\cdot)$ is the indicator function, $1 - \alpha > 0$ is the relative risk aversion, $\delta \leq 1$ is the disappointment threshold, and $\theta \geq 0$ is disappointment aversion. GDA preferences enable one to control the disappointment threshold by changing δ . Routledge and Zin's preferences nest two specifications. The expected utility is obtained by setting $\theta = 0$. Settings $\theta \neq 0$ and $\delta = 1$ reduce GDA preferences to the disappointment aversion utility.

B. Endowments and Inference Problem

I consider a Markov switching model for aggregate consumption growth

$$\Delta c_{t+1} = \mu_{s_{t+1}} + \sigma \varepsilon_{t+1}, \quad \varepsilon_{t+1} \sim N(0, 1),$$

where Δc_{t+1} is log consumption growth, s_{t+1} is a hidden two-state Markov chain with a state space $\mathcal{S} = \{1, 2\}$ and a transition matrix $\mathcal{P} = (\pi_{ij})$, in which $\pi_{11} = 1 - \pi_{12}$

and $\pi_{22} = 1 - \pi_{21}$ are transition probabilities, $\mu_{s_{t+1}}$ is the state-dependent mean growth rate, and σ is the constant consumption volatility. I assume $\mu_2 < \mu_1$ to identify $s_{t+1} = 1$ and $s_{t+1} = 2$ as expansion and recession, respectively.¹⁰

The motivation for constructing a two-state model is twofold. First, I want to maintain parsimony for the sake of convenient interpretation. Second, I do not introduce additional risks to isolate the impact of learning and GDA preferences. A model with additional ingredients would certainly make the framework more flexible. However, I show that a tightly calibrated GDA model with a single state variable can already reproduce the variance term structure with a wide array of salient features of the equity and derivatives markets.

I seek to price a levered consumption claim with log dividend growth

$$\Delta d_{t+1} = g_d + \lambda \Delta c_{t+1} + \sigma_d e_{t+1}, \quad e_{t+1} \sim N(0, 1),$$

in which λ is a leverage ratio on expected consumption growth. I use g_d to equalize long-run dividend and consumption growth rates, and σ_d to match the empirical dividend growth volatility. In addition, the chosen value of λ allows me to match the observed correlation between annual consumption and dividend growth rates.

The investor knows the true parameters and distribution of shocks but does not observe the state. At time t , the agent updates the probability of expansion $\pi_t = \mathbb{P}(s_{t+1} = 1 | \mathcal{F}_t)$ conditional on the history of consumption growth rates denoted by \mathcal{F}_t . I assume a Bayesian agent who updates his belief through Bayes' rule:

$$(3) \quad \pi_{t+1} = \frac{\pi_{11} f(\Delta c_{t+1} | 1) \pi_t + (1 - \pi_{22}) f(\Delta c_{t+1} | 2) (1 - \pi_t)}{f(\Delta c_{t+1} | 1) \pi_t + f(\Delta c_{t+1} | 2) (1 - \pi_t)},$$

$$f(\Delta c_{t+1} | i) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(\Delta c_{t+1} - \mu_i)^2}{2\sigma^2}}, \quad i = 1, 2.$$

III. Equilibrium

A. Equilibrium and Pricing Kernel

Following Routledge and Zin (2010), I show (see the Supplementary Material) that the gross return $R_{i,t+1}$ on the i th traded asset satisfies the condition

$$(4) \quad \mathbb{E}_t[M_{t+1} R_{i,t+1}] = 1,$$

in which M_{t+1} is the stochastic discount factor (SDF) of the GDA economy defined as

¹⁰The application of a regime-switching framework is a popular paradigm in the asset pricing literature. These models are flexible to embed business cycle fluctuations (Cecchetti, Lam, and Mark (1990), Veronesi (1999), Ju and Miao (2012), Johannes et al. (2016), and Collin-Dufresne et al. (2016)), the ‘‘peso problem’’ in the mean (Rietz (1988), Barro (2006), Backus, Chernov, and Martin (2011), and Gabaix (2012)) or persistence (Gillman et al. (2015)), long-run risks (Bonomo et al. (2011), (2015)), and economic recoveries (Hasler and Márfe (2016)) in endowments.

$$(5) \quad M_{t+1} = \underbrace{\beta \left(\frac{C_{t+1}}{C_t} \right)^{\rho-1}}_{M_{t+1}^{CRRA}} \cdot \underbrace{\left(\frac{V_{t+1}}{\mathcal{R}_t(V_{t+1})} \right)^{\alpha-\rho}}_{M_{t+1}^{EZ}} \cdot \underbrace{\left(\frac{1 + \theta \mathbb{I}(V_{t+1} \leq \delta \mathcal{R}_t(V_{t+1}))}{1 + \theta \delta^\alpha \mathbb{E}_t[\mathbb{I}(V_{t+1} \leq \delta \mathcal{R}_t(V_{t+1}))]} \right)}_{M_{t+1}^{GDA}}.$$

The first component M_{t+1}^{CRRA} is the SDF of the power utility. The second multiplier M_{t+1}^{EZ} is the adjustment of Epstein–Zin preferences, which separate the coefficient of risk aversion and EIS. The third component M_{t+1}^{GDA} represents the GDA adjustment. When the agent’s utility is below a predefined fraction of the certainty equivalent, more weight is attached to the SDF, magnifying the counter-cyclical dynamics of the pricing kernel. For a better understanding of the key role of GDA, I shut down the Epstein–Zin adjustment in SDF for the models with (generalized) disappointment aversion by setting $\alpha = \rho$. Thus, the pricing kernel simplifies to

$$M_{t+1} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{\rho-1} \cdot \left(\frac{1 + \theta \mathbb{I}(V_{t+1} \leq \delta \mathcal{R}_t(V_{t+1}))}{1 + \theta \delta^\alpha \mathbb{E}_t[\mathbb{I}(V_{t+1} \leq \delta \mathcal{R}_t(V_{t+1}))]} \right).$$

B. Model Solution

The latest long-run risk models generate significant nonlinearities, which, coupled with the log-linearization of equilibrium quantities, can generate economically significant numerical errors (Pohl et al. (2018)). Hence, I solve the model numerically using global solution methods to accurately capture the nonlinear nature of the model under consideration. The model solution boils down to approximating the return on the wealth portfolio R_{t+1}^ω and the equity return $R_{e,t+1}$ implicitly defined by equation (4). Denoting the investor’s wealth and equity price by W_t and P_t^e , we obtain

$$R_{t+1}^\omega = \frac{W_{t+1}}{W_t - C_t} = \frac{\frac{W_{t+1}}{C_{t+1}}}{\frac{W_t}{C_t} - 1} \cdot e^{\Delta c_{t+1}} \quad \wedge \quad R_{t+1}^e = \frac{P_{t+1}^e + D_{t+1}}{P_t^e} = \frac{\frac{P_{t+1}^e}{D_{t+1}} + 1}{\frac{P_t^e}{D_t}} \cdot e^{\Delta d_{t+1}}.$$

I conjecture that $\frac{W_t}{C_t} = G(\pi_t)$ and $\frac{P_t^e}{D_t} = H(\pi_t)$ are functions of π_t . I substitute R_{t+1}^ω and R_{t+1}^e into equation (4) and apply the projection method (Judd (1992)) to approximate $G(\pi_t)$ and $H(\pi_t)$. I discuss the numerical solution and its accuracy and provide the model-generated asset prices in the Supplementary Material.

IV. Data and Quantitative Results

A. Data

I construct annual real per capita consumption growth from Jan. 1930 to Dec. 2016 using the U.S. National Income and Product Accounts. I then retrieve data from the Center for Research in Security Prices to obtain aggregate equity market dividends and asset returns. To discipline quantitative analysis, I tightly calibrate each model in this article to closely match the key moments of fundamentals and equity returns.

In addition to standard asset pricing moments, I study the implications of different models for the high moment risk premiums and option prices. The variance premium is the difference between expectations of stock market return variance under the risk-neutral \mathbb{Q} and actual physical \mathbb{P} probability measures.¹¹ Formally, a τ -month variance premium at time t is $vp_t = \mathbb{E}_t^{\mathbb{Q}}[\text{RETURN_VARIATION}(t, t + \tau)] - \mathbb{E}_t^{\mathbb{P}}[\text{RETURN_VARIATION}(t, t + \tau)]$, in which the total return variation is calculated over the period t to $t + \tau$. The quantity vp_t corresponds to the expected profit of a variance swap, which pays the equity's realized variance over the term of the contract. Like the variance premium, I follow Kozhan, Neuberger, and Schneider (2013) and define a τ -month skew risk premium at time t as $sp_t = \frac{\mathbb{E}_t^{\mathbb{P}}[\text{RETURN_SKEWNESS}(t, t + \tau)]}{\mathbb{E}_t^{\mathbb{Q}}[\text{RETURN_SKEWNESS}(t, t + \tau)]} - 1$, in which the total return skewness is calculated from t to $t + \tau$. The quantity sp_t corresponds to the excess return on a skew swap, which pays the equity's realized skewness over the term of the contract. The literature has mainly focused on the variance premium, while the skew premium has received little attention, especially from theoretical research.

The data for the variance premium covers the period from Jan. 1990 to Dec. 2016 and is from the Chicago Board Options Exchange (CBOE). For the skew risk premium and implied volatility surface, I use European options written on the S&P 500 index and traded on the CBOE. The options data cover the period from Jan. 1996 to Dec. 2016 and are from OptionMetrics.¹² Table 1 shows summary statistics for 1-month variance and skew risk premiums.¹³ Figure 2 shows the implied volatility curves. The size of the variance and skew premiums as well as the level and the slope of implied volatility curves remain a challenge for asset pricing models. This article shows that a model with GDA preferences and learning about rare depressions jointly captures standard moments of equity returns, high moment premiums, and option prices with new evidence about the variance term structure.

TABLE 1
Summary Statistics: Variance and Skew Risk Premiums

Table 1 reports monthly descriptive statistics for the conditional variance vp_t and skew sp_t premiums. Mean, median, Std. Dev., max, skewness, and kurtosis report the sample average, median, standard deviation, maximum, skewness, and kurtosis, respectively. The empirical statistics of the variance and skew risk premiums are for the U.S. data from Jan. 1990 to Dec. 2016 and from Jan. 1996 to Dec. 2016, respectively.

	vp_t	sp_t
Mean	10.24	-42.12
Median	7.50	-68.11
Std. Dev.	10.49	82.11
Max	83.70	447.37
Skewness	2.62	3.57
Kurtosis	14.15	16.26

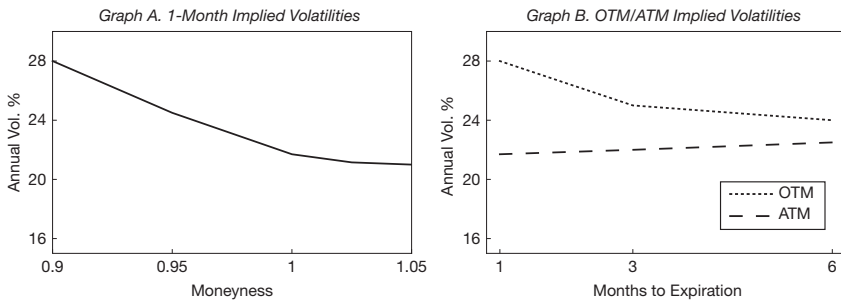
¹¹In the model, the Radon–Nikodym derivative is defined as $\frac{d\mathbb{Q}}{d\mathbb{P}} = \frac{M_{t+1}}{\mathbb{E}_t(M_{t+1})}$ and allows one to compute the risk-neutral moments.

¹²I present the empirical methodology and the model-based asset prices in Supplementary Material.

¹³The estimates are consistent with Bakshi, Kapadia, and Madan (2003), Bollerslev et al. (2009), and Kozhan et al. (2013).

FIGURE 2
Implied Volatilities

Graph A of Figure 2 plots the empirical 1-month implied volatility curve as a function of moneyness. Graph B plots the empirical implied volatility curves for ATM and OTM options as functions of the time to maturity (in months). All curves are for the U.S. data from Jan. 1996 to Dec. 2016.



B. Calibration

To better understand the role of GDA, I consider three frameworks: a model with GDA preferences (GDA), an economy with disappointment aversion preferences (DA), and a specification with Epstein–Zin preferences (EZ). The comparison of GDA and DA isolates the contribution of disappointment aversion, whereas the comparison of GDA and EZ illustrates the impact of the agent’s preference for early resolution of uncertainty. Having solved the model numerically, I generate 10,000 simulations of each calibration and report model-based 5th, 50th, and 95th percentiles of sample moments of cash flows and asset prices across all simulations.¹⁴ In line with the data, the model-implied cash flows and returns are based on simulations with depressions, while the model-based variance forwards, moment risk premiums, and option prices correspond to simulations without depressions. The results are robust to the inclusion of rare events, which are excluded to eliminate the impact of large consumption declines and to highlight the role of learning and GDA.

Table 2 reports the parameter values. As in Bansal and Yaron (2004), I make the model’s time-averaged consumption statistics consistent with observed annual log consumption growth. As in Collin-Dufresne et al. (2016), I calibrate the recession state to a consumption decline in the United States during the Great Depression.¹⁵ Specifically, I set $\pi_{11} = 1,151/1,152$ and $\pi_{22} = 47/48$. These numbers imply an average duration of the high-growth state of $(1 - \pi_{11})^{-1} = 96$ years

¹⁴The previous version of the paper reported model population moments. For a convenient exposition of tables and figures, those results are not reported but are available from the author. In those results, I check that the fact the model explains the variance term structure and other moments is not a finite-sample phenomenon.

¹⁵The Great Depression is the only example of a consumption disaster in U.S. history for the period considered in my paper. Thus, I naturally calibrate the recession state to this observation following Collin-Dufresne et al. (2016). Furthermore, Nakamura, Steinsson, Barro, and Ursua (2013) note that rare disasters tend to unfold over multiple years. Instead of assuming extreme instantaneous consumption disasters, I choose a milder depression with an average duration corresponding to 4 years of the Great Depression.

TABLE 2
Parameter Values

Table 2 reports parameter values in the cash-flow processes and the three models: GDA, DA, and EZ.

Parameter	Description	Value		
π_{11}	Transition probability from expansion to expansion	1, 151/1, 152		
π_{22}	Transition probability from recession to recession	47/48		
$\mu_1 \times 12$	Consumption growth in expansion	2.06		
$\mu_2 \times 12$	Consumption growth in recession	-4.6		
$g_d \times 12$	Mean adjustment of dividend growth	-2.87		
$\sigma \times \sqrt{12}$	Std. Dev. of consumption growth shock	2.6		
$\sigma_d \times \sqrt{12}$	Std. Dev. of dividend growth shock	11.41		
λ	Leverage ratio	2.6		
		GDA	DA	EZ
β^{12}	Discount factor	0.99	0.99	0.99
$1/(1-\rho)$	EIS	1.5	1.5	1.5
$1-\alpha$	Risk aversion	1/1.5	1/1.5	6.0
θ	Disappointment aversion	8.41	0.6	0
δ	Disappointment threshold	0.930	0	

and the depression state of $(1 - \pi_{22})^{-1} = 4$ years. The unconditional probability of expansion is $\bar{\pi}_{11} = (1 - \pi_{22}) / (2 - \pi_{11} - \pi_{22}) = 0.96$ and hence the economy experiences one 4-year depression per century, consistent with the historical data. Consumption declines on average at the annual rates of $\mu_2 \times 12 = -4.6\%$ in the depression state, which is equal to an average annual decline in the real, per capita log consumption growth during the Great Depression.

I now calibrate parameters in the dividend process. To compare my results to prior studies, particularly the disaster literature, I set the leverage ratio $\lambda = 2.6$, the value used in Seo and Wachter (2019).¹⁶ I further follow the literature and set g_d to equalize the long-run dividend and consumption growth. The standard deviation of the dividend process σ_d is used to generate large annual dividend volatility observed in the data.

Table 2 further summarizes the values of GDA, DA, and EZ preferences. I set $\beta^{12} = 0.99$ and $1/(1-\rho) = 1.5$ in all cases. In the GDA model, the coefficient of relative risk aversion is $1 - \alpha = 1/1.5$. This cancels the Epstein–Zin adjustment in SDF as shown in Section III and also deletes one degree of freedom caused by extra GDA parameters. I jointly set $\theta = 8.41$ and $\delta = 0.930$ to match the high equity premium. The calibrated disappointment aversion is consistent with the empirical estimates from 3.29 to 8.41 (Delikouras (2017)). Note that the variance term structure, the variance and skew premiums, and the implied volatility surface are not directly targeted during the model calibration.

In the DA model, I set $1 - \alpha = 1 - \rho = 1/1.5$ to eliminate the impact of a relative risk aversion parameter on SDF. I also shut down the GDA channel by setting $\delta = 1$. This inevitably generates larger effective risk aversion in good times due to an increased number of disappointing events, significantly distorting equity

¹⁶I regress the annual dividends on the annual consumption covering the period 1930–2016 and find the leverage ratio is around 2.5, a number within an interval of commonly used values from 1.5 to 4.5. The leverage ratio is an important parameter for two reasons. First, it controls the volatility of dividends in normal times. Second, it determines the decline of dividends in the depression state. Consequently, a larger leverage parameter would increase the payoff of put options, conditional on the depression realization.

moments in the DA model. Thus, I decrease the disappointment aversion parameter $\theta = 0.6$ to match the observed equity premium. The remaining parameters are fixed at the initial values. For the EZ model, I turn off disappointment aversion by setting $\theta = 0$. The model operates only through the risk aversion channel with the coefficient of relative risk aversion of $1 - \alpha = 6$. In this case, the agent has a preference for early resolution of uncertainty, a workhorse in the asset pricing literature. Other parameters correspond to those in the GDA model.

C. Endowments and Equity Returns

Panel A of Table 3 compares the annualized consumption and dividends moments of the data with those implied by the calibration. A two-state regime-switching process matches the key empirical statistics well. Panel B of Table 3 reports the annualized moments of equity returns for the three specifications. All three models do a good job of accounting for salient features of equity returns, as all predict the low risk-free rate, the large equity premium, and the volatility of excess returns. Also, the volatility of the risk-free rate and the level of the log price-dividend ratio correspond well to the empirical estimates. The shortcoming of the three models is too low volatility of the log price-dividend ratio.

D. The Price of Variance Risk

Figure 3 compares the empirical and model-based term structure of variance swap prices and returns. Graph A shows that the GDA model does a good job of matching the overall shape of annualized Sharpe ratios. In particular, it generates a curve that is negative and steep at shorter horizons and becomes positive and

TABLE 3
Cash Flows and Stock Market Returns

Panel A of Table 3 reports moments of consumption and dividend growth denoted by Δc and Δd . Panel B reports moments of the log risk-free rate r_f , the excess log equity returns $r_e - r_f$, and the log price-dividend ratio PD. The entries are annualized statistics except for autocorrelation and correlation. The moments are for the data and the three models: GDA, DA, and EZ. The empirical moments are for the U.S. data from Jan. 1930 to Dec. 2016. For each model, I simulate 10,000 economies at a monthly frequency with a sample size equal to its empirical counterpart and report percentiles of sample statistics. The model-implied results are based on the simulations with consumption disasters, consistent with the historical data. I use common notations for mean E , volatility σ , autocorrelation AC1, and correlation Corr.

	Data	GDA			DA			EZ		
		5%	50%	95%	5%	50%	95%	5%	50%	95%
<i>Panel A. Cash Flows</i>										
$E(\Delta c)$	1.83	0.91	1.85	2.40	0.91	1.85	2.40	0.91	1.85	2.40
$\sigma(\Delta c)$	2.22	1.90	2.28	3.19	1.90	2.28	3.19	1.90	2.28	3.19
AC1(Δc)	0.50	0.09	0.30	0.62	0.09	0.30	0.62	0.09	0.30	0.62
$E(\Delta d)$	1.44	-1.10	1.91	4.44	-1.10	1.91	4.44	-1.10	1.91	4.44
$\sigma(\Delta d)$	11.04	9.51	11.05	12.97	9.51	11.05	12.97	9.51	11.05	12.97
AC1(Δd)	0.19	0.09	0.27	0.46	0.09	0.27	0.46	0.09	0.27	0.46
Corr($\Delta c, \Delta d$)	0.55	0.38	0.55	0.71	0.38	0.55	0.71	0.38	0.55	0.71
<i>Panel B. Returns</i>										
$E(r_f)$	0.81	-0.13	0.86	1.49	0.68	1.14	1.20	0.22	1.03	1.41
$\sigma(r_f)$	1.87	1.48	2.52	3.51	0.04	0.25	1.22	0.73	1.50	2.34
$E(r_e - r_f)$	5.22	3.67	6.10	8.35	3.43	6.04	8.47	3.50	5.89	8.19
$\sigma(r_e - r_f)$	19.77	15.58	19.22	23.11	13.03	16.02	20.34	14.64	18.69	23.49
$E(PD)$	3.11	2.96	3.03	3.05	2.90	2.97	2.98	2.95	3.04	3.06
$\sigma(PD)$	0.33	0.04	0.08	0.18	0.01	0.05	0.18	0.03	0.08	0.22

FIGURE 3

Sharpe Ratios and Forward Variance Claim Prices

Figure 3 plots annualized Sharpe ratios and average prices for forward variance claims for the data and the three models: GDA, DA, and EZ. The prices are reported in annualized volatility terms. The empirical lines are from Dew-Becker et al. (2017) and correspond to the U.S. data from 1996 to 2013. For each model, I simulate 10,000 economies at a monthly frequency with a sample size equal to its empirical counterpart and report medians of sample statistics. The model-implied results are based on the simulations without consumption disasters, consistent with the historical data.

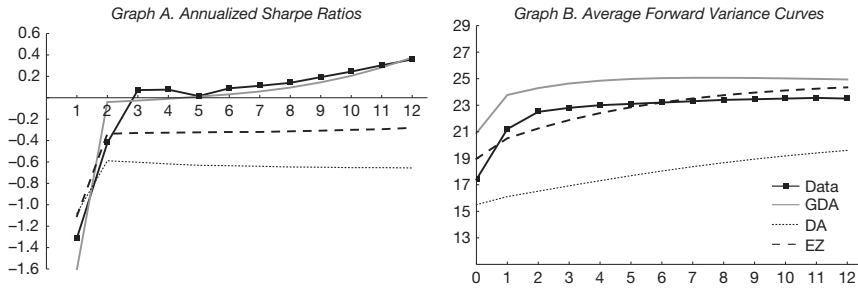


TABLE 4

Model Tests Using Annualized Sharpe Ratios for Forward Variance Claims

The entries in Table 4 are for the three models: GDA, DA, and EZ. For each model, I simulate 10,000 economies at a monthly frequency with a sample size equal to the length of the variance swap data. In each simulation, I calculate average annualized Sharpe ratios for forward variance claims with 1-, 3-, and 12-month maturities. For each model, the first row shows fractions of samples in which the simulated Sharpe ratios are at least as small as the empirical 1-month estimates. The second and third rows present the fraction of samples in which the simulated Sharpe ratios are at least as large as the empirical 3-month and 12-month estimates, respectively. The entries of the bottom row are the fraction of samples in which all three conditions are satisfied simultaneously. The model-implied results are based on the simulations without consumption disasters, consistent with the historical data.

	p-Value		
	GDA	DA	EZ
Simulated 1mo/SR \leq empirical SR	0.91	0.26	0.27
Simulated 3mo/SR \geq empirical SR	0.38	<0.01	0.03
Simulated 12mo/SR \geq empirical SR	0.48	<0.01	<0.01
Joint test: 1mo/SR \leq data \wedge 3mo/SR \geq data \wedge 12mo/SR \geq data	0.32	<0.01	<0.01

upward-sloping at longer maturities. The figure also shows that both DA and EZ specifications fail to reconcile the concave and upward shape of the term structure. Consistent with Dew-Becker et al. (2017), the calibration with Epstein–Zin preferences underprices variance risk in the short term and overprices future variance in the long term. The DA model implies even more negative Sharpe ratios at longer horizons, while the 1-month forwards earn a similar risk premium as in the EZ model. Graph B plots the average prices of forward variance claims for different maturities in the data and the three models. The empirical curve is steep and concave at the very short end and it flattens significantly at the long end. In contrast, the DA and EZ specifications predict strongly upward-sloping term structures at all horizons. Although the GDA model generates slightly higher prices of variance claims, it captures the concave shape and the flatness of the curve at longer maturities.

Table 4 augments the results in Figure 3 by reporting the p -values of annualized Sharpe ratios with respect to their finite-sample distribution. For each model, it shows the fraction of samples across 10,000 simulations of the economy satisfying

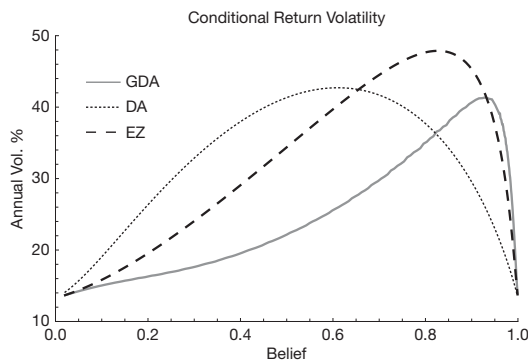
one of the conditions. For the first three conditions, simulated average Sharpe ratios for 1-, 3-, and 12-month horizons should be, respectively, smaller, larger, and larger than the empirical estimates. One can interpret these fractions as p -values for a one-sided test of the model generating as negative or as positive average Sharpe ratios for a particular maturity as in the data. For the last condition, simulated statistics should jointly satisfy the first three requirements. This corresponds to the p -value for a test of the model replicating the observed upward-sloping shape of the term structure.

Table 4 shows that we cannot reject any of the three models based on the 1-month variance forward returns only. Specifically, one would expect to see as small average 1-month Sharpe ratios as observed empirically in 91%, 26%, and 27% of the time in the GDA, DA, and EZ specifications, respectively. At longer maturities, however, one can reject at the 5% level the null hypothesis that the DA or EZ frameworks generate the variance swap data. The GDA model instead generates large p -values for all tests and cannot be rejected. In particular, the models with disappointment aversion or Epstein–Zin preferences would predict positive Sharpe ratios at longer maturities as in the data in fewer than 3% of simulations, while the likelihood of replicating the overall shape is less than 1%. This is in stark contrast to the GDA model, which captures negative Sharpe ratios at the short end and positive ones at the long end in 32% of the simulations.

To gain a better understanding of the results, Figure 4 illustrates annualized return volatility as a function of the posterior probability of the expansion in the three models. The volatility has a pronounced humped shape and is maximized at an interior point of the probability simplex in all cases. The GDA model generates highly skewed conditional volatility. The DA specification yields a symmetric shape of the volatility curve. The volatility line in the EZ economy is roughly located in the middle of the two. For Epstein–Zin and especially disappointment aversion preferences, return volatility is high for a wide range of beliefs and becomes low only when the investor has full confidence in the state. As the

FIGURE 4
Return Volatility

Figure 4 plots equity return volatility as a function of a posterior belief for the three models: GDA, DA, and EZ. Quantities are reported in annualized volatility terms, $100 \times \sqrt{12 \times \text{var}_t(r_e)}$.



investor's beliefs tend to change slowly over time, high return volatility persists in the long term and hence increases the hedge against long-term volatility risk. As a result, this generates the upward-sloping term structure of variance claim prices, which is inconsistent with the data.¹⁷

In contrast, return volatility in the GDA model is high within a narrow range of beliefs and quickly diminishes outside this interval. When the investor's beliefs become pessimistic, return volatility initially spikes but does not persist in the long term, implying a larger amount of variance risk in the short term. In equilibrium, the properties of return variance transmit to variance forwards, which generates the inversion in their prices in bad times. In the short term, the inversion increases the hedge against realized variance. In the long run, it kills the upward-sloping effect of time-varying disaster risk and produces the flat unconditional term structure of prices. I also show that the inversion is strong enough to produce on average positive and slightly increasing Sharpe ratios at longer maturities.

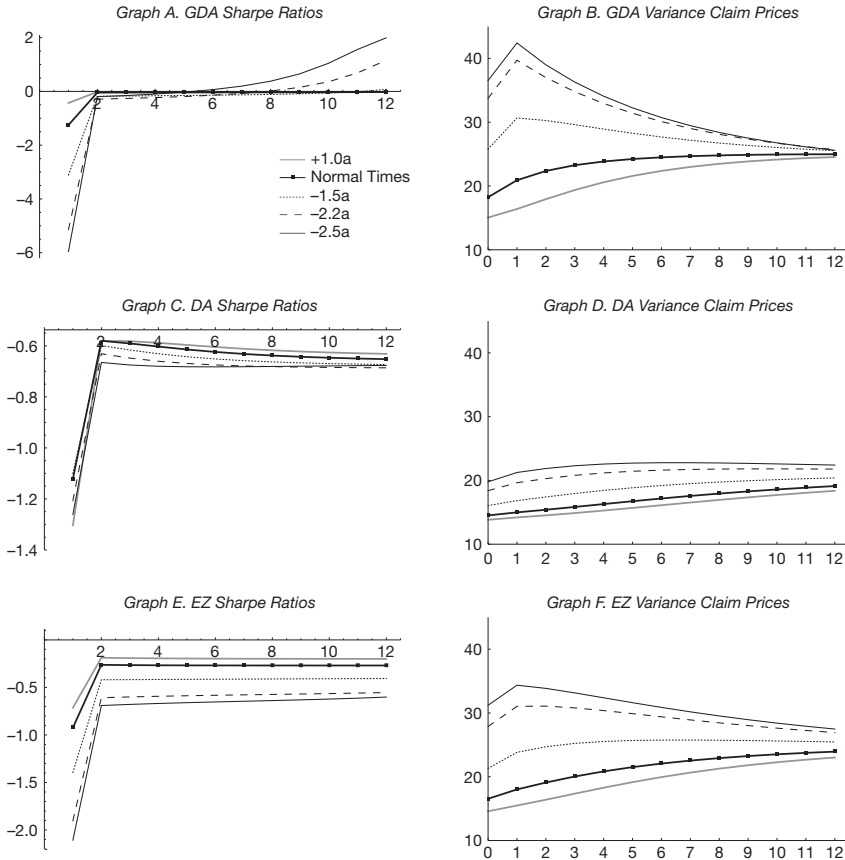
The mechanism determining the conditional return volatility is as follows: Veronesi (1999) demonstrates that, in the endowment economy with two hidden regimes, price sensitivity to news is strongly driven by the risk aversion component stemming from the investor's degree of risk aversion. In the DA model, risk aversion is equally large in the expansion and depression states because a high disappointment threshold implies a large number of disappointing outcomes in the two regimes. Thus, price sensitivities are similar across states, resulting in symmetric conditional return volatility. In the EZ economy, equity prices are more sensitive to consumption shocks in good times than in bad times, although Epstein–Zin preferences do not generate a significant difference in price sensitivities in the two regimes. In the GDA model, instead, substantial countercyclical risk aversion leads to a stronger overreaction of stock prices to bad news in good times, whereas equity prices are substantially less sensitive to good news in bad times. This asymmetry in price sensitivities leads to strongly skewed return volatility in the GDA model.

As an additional exercise, Figure 5 provides impulse responses of the conditional term structure of Sharpe ratios and average prices. The investor holds a median belief (normal times). I then study conditional dynamics of the term structures next period when consumption growth is 1.5 standard deviations above and 1.5, 2.2, and 2.5 standard deviations below average growth in the expansion (good and bad times). Figure 5 shows that the DA and EZ models predict negative Sharpe ratios for all economic conditions. Contrary to the empirical evidence, average Sharpe ratios become more negative in the upside scenario under disappointment aversion. The economy with GDA preferences generates a procyclical and steep curve for short-term claims, consistent with Ait-Sahalia et al. (2020). Furthermore, the term structure of Sharpe ratios is insignificant for maturities longer than 2 months in good and normal times as well as bad but not depression-like states and is steep and positive in response to large consumption declines.

¹⁷This intuition also applies to the models with full information. If the state is observable, the high stock market variance happens during consumption disasters and hence persists on average for 4 years, the average duration of a depression state. As a result, high variance risk will concentrate in the long term, and forward variance will be upward-sloping, counter to what we observe empirically.

FIGURE 5
Conditional Sharpe Ratios and Forward Variance Claim Prices

Figure 5 plots annualized Sharpe ratios and average prices for forward variance claims for the three models: GDA, DA, and EZ. Each graph shows the term structures in good, normal, and bad times. The economy is initially in normal times, corresponding to a median posterior belief. In good (bad) times, denoted by $+1.0\sigma$ (-1.5σ , -2.2σ , and -2.5σ), consumption growth is 1.0 (1.5, 2.2, and 2.5) standard deviation(s) above (below) an average growth in expansion.



The latter feature of the GDA model enables to match the sign and shape of the unconditional curves.

Figure 5 further depicts impulse responses of variance claim prices. The average curve for the DA model remains upward-sloping in all scenarios. This explains negative average returns on variance forwards. For the EZ economy, the term structure of prices switches from strongly increasing in normal and good times to slightly increasing in bad (but not severe) times, and it even becomes weakly downward-sloping in very bad times. Nevertheless, this amplification of short-term prices is too weak to generate on average positive returns on holding a variance forward. In contrast, GDA inverts the term structure in all bad scenarios, and this inversion is substantially stronger than in the EZ economy. Thus, GDA preferences strongly amplify the short-term variance risk in bad times that enables one to replicate empirical term structures.

Next, I conduct a sensitivity analysis to examine the robustness of key results to alternative calibrations of preference specifications and to address the concern that the findings are driven by a particular choice of parameters. Specifically, I change one key parameter in each of the three preference specifications, while holding the remaining parameters as in the original calibration. In the GDA model, I consider smaller and larger values of disappointment aversion and threshold parameters. In the DA model, I decrease or increase disappointment aversion compared to the original calibration. In the EZ model, I consider smaller and larger relative risk aversion coefficients.

Figure 6 depicts Sharpe ratios and prices for variance forwards in various calibrations of GDA preferences. The shape of variance forward prices flattens and the term structure of Sharpe ratios becomes upward-sloping with the higher disappointment threshold or disappointment aversion. Intuitively, variance risk is amplified more in the short term than in the long term in bad times. As a result, this generates downward- and upward-sloping patterns in prices and Sharpe ratios, respectively. In normal and good times, average prices are slightly increasing in the horizon. However, only short-term variance risk earns a significant premium as measured by large and negative Sharpe ratios for 1 and 2 months but insignificant ratios for longer horizons. Since higher disappointment risk reinforces the first effect, the higher disappointment threshold or disappointment aversion implies flatter and steeper term structures of prices and Sharpe ratios, respectively.

FIGURE 6
Sensitivity of Sharpe Ratios and Forward Variance Claim Prices: GDA

Figure 6 plots annualized Sharpe ratios and average prices for forward variance claims for different model calibrations with generalized disappointment aversion preferences. GDA corresponds to the original GDA model. In GDA_{θ_1} and GDA_{θ_2} , $\theta_1 = 6.41$ and $\theta_2 = 10.41$. In GDA_{δ_1} and GDA_{δ_2} , $\delta_1 = 0.920$ and $\delta_2 = 0.940$. If not stated otherwise, the remaining parameters are set at the original values in the GDA model. For each model, I simulate 10,000 economies at a monthly frequency with a sample size equal to its empirical counterpart and report medians of sample statistics. The model-implied results are based on the simulations without consumption disasters, consistent with the historical data.

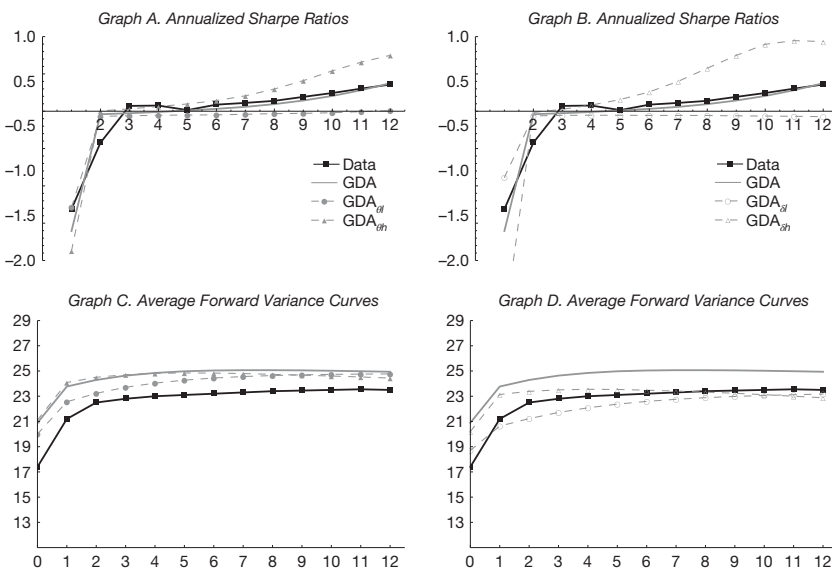


FIGURE 7

Sensitivity of Sharpe Ratios and Forward Variance Claim Prices: DA and EZ

Figure 7 plots annualized Sharpe ratios and average prices for forward variance claims for different model calibrations with disappointment aversion and Epstein–Zin preferences. DA and EZ correspond to the original DA and EZ models. In DA_{θ_1} and DA_{θ_2} , $\theta_1 = 0.5$ and $\theta_2 = 0.7$. In $EZ_{(1-\alpha)_1}$ and $EZ_{(1-\alpha)_2}$, $(1-\alpha)_1 = 5$ and $(1-\alpha)_2 = 7$. If not stated otherwise, the remaining parameters are set at the original values in the DA and EZ models. For each model, I simulate 10,000 economies at a monthly frequency with a sample size equal to its empirical counterpart and report medians of sample statistics. The model-implied results are based on simulations without consumption disasters, consistent with the historical data.

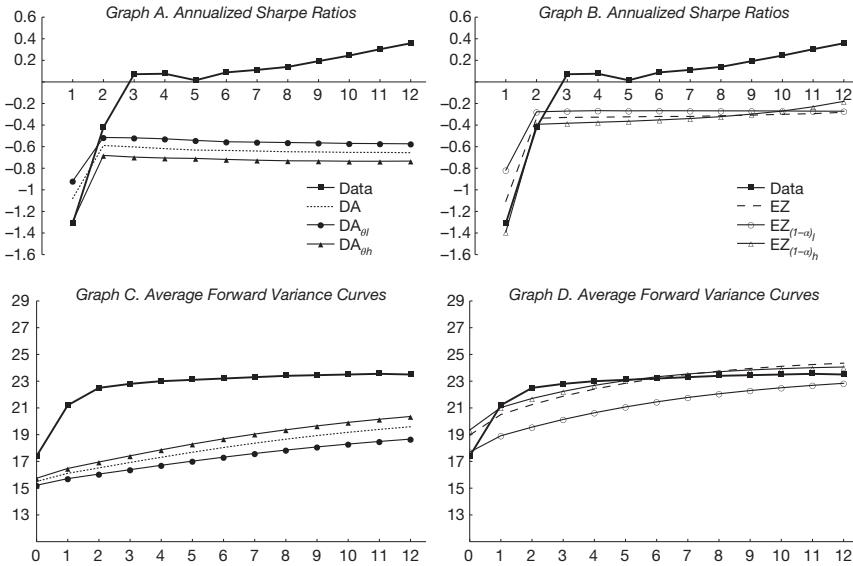


Figure 7 examines the impact of disappointment and risk aversion parameters on the variance term structures in the models with Gul and Epstein–Zin preferences. In the DA specification, the slope of forward variance prices is increasing in disappointment aversion. The reason is that the disappointment-averse investor strongly dislikes low and high variance. Thus, stronger disappointment aversion increases already high insurance premia against shocks to realized and future volatility. In the EZ economy, the slope of forward variance prices is decreasing in risk aversion. To generate a close-to-zero slope at least after the 10-month maturity, the risk aversion should be at least 7. For this value, however, the model would generate a Sharpe ratio of less than -2.0 for the 1-month claim compared to -1.3 in the data. Moreover, with this value of relative risk aversion, the mean equity premium has a median value of 8% in the EZ model, well above the empirical estimate of around 5%. Raising risk aversion even more would only worsen the model fit with the variance term structure at the 1-month maturity and with equity moments and higher-moment risk premiums (see the Supplementary Material). Thus, one cannot reconcile the variance term structure in the EZ framework by increasing risk aversion.

In sum, this sensitivity analysis confirms that the pricing kernel, necessary to reconcile the empirical variance term structure, is consistent with GDA and cannot be supported by parameter values in alternative preferences. Supplementary

Material augments a comparative statics exercise by reporting the remaining results in alternative model calibrations. It demonstrates that the DA and EZ specifications with different parameter choices are unable to capture the higher-moment risk premiums and implied volatility curves.

E. Variance and Skew Risk Premiums

Panel A of Table 5 collects moments of the variance premium and related measures in the data and models. The GDA economy is able to generate a large and volatile variance premium. It also qualitatively respects the non-normality of the variance premium distribution, although the sample skewness and kurtosis statistics are smaller relative to the data. The GDA model accounts for the variance premium with empirically consistent conditional return variances under both probability measures. In particular, it predicts that return variance is more volatile under the risk-neutral distribution and that both variances are persistent, as they are in the data.

I now examine return predictability by the variance premium documented by prior literature. I regress the 1-, 3-, and 6-month cumulative excess log returns (expressed in percentages) on the lagged monthly variance premium. Panel B of Table 5 reports positive and slightly decreasing regression coefficients and

TABLE 5
Variance Premium and Predictability

Panel A of Table 5 reports moments of variance premium vp , market return variances $\text{var}_t^{\mathbb{P}}(r_e)$ and $\text{var}_t^{\mathbb{Q}}(r_e)$ under physical \mathbb{P} and risk-neutral \mathbb{Q} probability measures. The Panel A entries are monthly statistics. Panel B reports results of the predictive regression of h -month future excess log equity returns constructed as $r_{t+1 \rightarrow t+h}^{\text{ex}} = \sum_{i=1}^h (r_{e,t+i} - r_{f,t+i})$ on the lagged variance premium VP_t . Specifically, the slope estimates $\beta(h)$ and $R^2(h)$ are based on the linear projection:

$$100 \times r_{t+1 \rightarrow t+h}^{\text{ex}} = \text{INTERCEPT} + \beta(h) \times VP_t + \varepsilon_{t+h}, h = 1, 3, 6.$$

The moments and regression outputs are for the data and the three models: GDA, DA, and EZ. The empirical statistics are for the U.S. data from Jan. 1990 to Dec. 2016. For each model, I simulate 10,000 economies at a monthly frequency with a sample size equal to its empirical counterpart and report percentiles of sample statistics based on these series. The model-implied results are based on the simulations without consumption disasters, consistent with the historical data. I use common notations for mean E , volatility σ , autocorrelation AC1, skewness SKEW, and kurtosis KURT.

	Data	GDA			DA			EZ		
		5%	50%	95%	5%	50%	95%	5%	50%	95%
<i>Panel A. Variance Premium</i>										
$E(VP)$	10.27	8.13	12.32	17.14	1.34	2.11	3.39	3.15	4.92	7.23
$\sigma(VP)$	10.87	12.22	15.99	18.93	1.53	3.11	5.25	4.79	7.46	10.91
SKEW(VP)	2.33	0.91	1.49	2.25	0.49	2.76	4.17	-0.62	1.71	2.98
KURT(VP)	10.90	2.34	4.06	7.67	5.88	12.03	24.96	4.00	7.27	13.59
$\sigma(\text{var}_t^{\mathbb{P}}(r_e))$	29.32	17.34	25.44	32.68	5.02	14.25	36.48	13.00	25.50	40.47
$AC1(\text{var}_t^{\mathbb{P}}(r_e))$	0.79	0.70	0.81	0.88	0.61	0.79	0.92	0.66	0.82	0.91
$\sigma(\text{var}_t^{\mathbb{Q}}(r_e))$	33.76	29.58	40.60	49.11	6.55	17.24	38.91	17.79	31.46	45.00
$AC1(\text{var}_t^{\mathbb{Q}}(r_e))$	0.80	0.69	0.79	0.85	0.62	0.80	0.92	0.66	0.82	0.89
SKEW($\text{var}_t^{\mathbb{P}}(r_e)$)	3.53	0.86	1.47	2.21	2.40	3.73	5.75	1.47	2.30	3.34
KURT($\text{var}_t^{\mathbb{P}}(r_e)$)	21.47	2.30	4.13	7.78	8.46	19.47	45.54	3.87	8.24	16.18
<i>Panel B. Predictability of Excess Returns</i>										
$\beta(1m)$	0.76	0.19	0.75	1.38	-1.43	1.81	5.56	-0.95	0.93	2.77
$R^2(1m)$	2.70	0.15	2.39	6.99	0.02	1.09	4.19	0.01	1.31	6.01
$\beta(3m)$	0.83	0.18	0.63	1.09	-1.31	1.55	4.02	-0.73	0.87	2.11
$R^2(3m)$	8.61	0.47	5.57	15.24	0.02	2.39	8.67	0.03	3.28	12.87
$\beta(6m)$	0.57	0.15	0.50	0.82	-1.05	1.24	3.06	-0.63	0.74	1.60
$R^2(6m)$	7.55	0.68	7.78	20.79	0.08	3.32	13.14	0.04	4.67	16.95

increasing R^2 s with the horizon. The GDA model matches the magnitude and monotonicity of coefficients and R^2 statistics.

Table 5 shows that the model with disappointment aversion preferences produces a mean and volatility of the variance premium that are more than five times smaller than with the generalized utility function. Turning off the GDA channel also leads to a significant reduction in the volatility of return variance in the DA model. As the variance premium decreases, its predictive power for the excess log returns also suffers. This is manifested in the lower R^2 s and empirically inconsistent regression coefficients. Next, I turn off any source of (generalized) disappointment aversion and consider a representative agent with Epstein–Zin preferences. The EZ model leads to around a twofold increase in the mean and volatility of the variance premium relative to the DA model, but sample statistics are less than half of the numbers in the GDA model. A smaller variance premium is due to the reduced volatility in conditional variances. A smaller variance premium in the DA and EZ models results in excessively high regression coefficients and too small R^2 s in the predictive regressions.

Table 6 reports summary statistics of the skew risk premium in the data and models. The GDA model produces a sizeable skew premium, which corresponds well to the historical value and generates positive skewness and excess kurtosis statistics. The conditional mean of the return skewness under both measures is significantly negative, although the model cannot fully capture the size observed in the data. The main drawback of the GDA model is lower volatility of the skew premium, realized and implied skew. Since conditional dynamics of the model are driven by a single state, allowing the model to operate through other channels (time-varying expected growth and volatility, jumps in consumption, etc.) would make the economy more flexible to jointly match all moments.

Table 6 also shows that disappointment aversion predicts the wrong sign of the skew premium. The DA model also predicts the smallest first and second moments of return skewness across the three models. In the EZ model, the risk-neutral return density becomes more distorted toward the left tail; however, the model generates

TABLE 6
Skew Premium

Table 6 reports the moments of skew premium SP, market return skewness $SKEW_t^P(r_e)$ and $SKEW_t^Q(r_e)$ under the physical \mathbb{P} and risk-neutral \mathbb{Q} probability measures. The entries are monthly statistics. The moments are for the data and the three models: GDA, DA, and EZ. The empirical statistics are for the U.S. data from Jan. 1996 to Jan. 2016. For each model, I simulate 10,000 economies at a monthly frequency with a sample size equal to its empirical counterpart and report percentiles of sample statistics based on these series. The model-implied results are based on the simulations without consumption disasters, consistent with the historical data. I use common notations for mean E , volatility σ , skewness SKEW, and kurtosis KURT.

	Data	GDA			DA			EZ		
		5%	50%	95%	5%	50%	95%	5%	50%	95%
$E(SP)$	-42.20	-39.11	-34.58	-30.78	26.58	34.88	56.36	-22.79	-19.34	-12.84
$\sigma(SP)$	81.81	11.23	26.42	46.52	24.44	29.53	377.79	3.03	21.31	91.65
$SKEW(SP)$	3.57	-4.32	3.28	8.56	-3.37	1.24	13.98	-11.70	3.23	13.69
$KURT(SP)$	16.26	1.93	43.48	112.04	3.10	4.65	215.60	2.04	80.40	219.37
$AR1(SP)$	0.04	-0.12	0.15	0.62	-0.01	0.61	0.70	-0.27	0.11	0.58
$E(SKEW_t^P(r_e))$	-87.52	-42.55	-39.99	-33.83	-20.49	-15.82	-12.34	-38.00	-33.98	-29.43
$\sigma(SKEW_t^P(r_e))$	173.59	8.99	11.79	22.21	7.68	11.41	15.56	12.03	13.61	23.97
$E(SKEW_t^Q(r_e))$	-177.73	-70.44	-64.13	-53.83	-17.39	-13.22	-9.82	-47.60	-42.42	-36.45
$\sigma(SKEW_t^Q(r_e))$	92.33	23.20	28.27	41.96	7.91	11.14	15.54	16.44	18.68	28.90

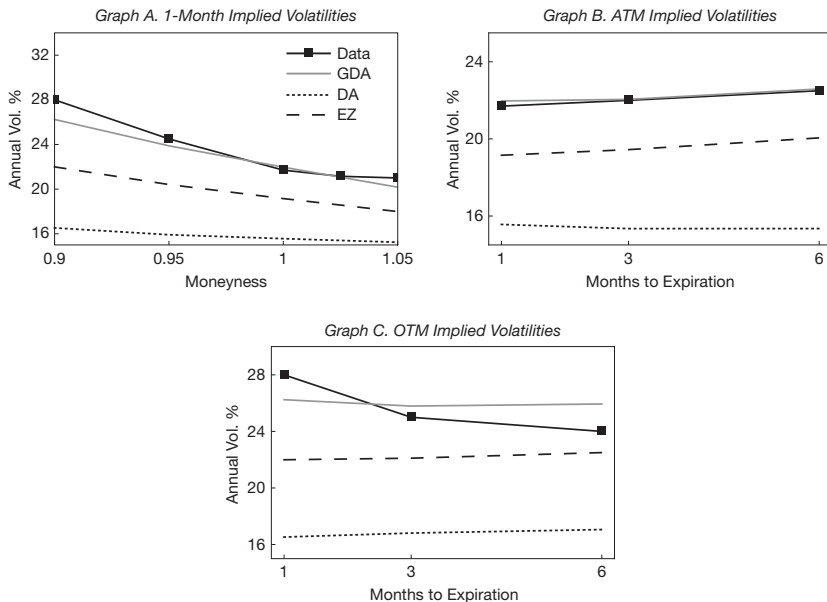
less than half of the average skew premium in the data. Although the EZ model predicts the correct sign, it significantly understates the magnitude. Overall, GDA better explains salient features of the skew premium than nested preferences.

F. The Term Structure of Implied Volatilities

Figure 8 compares the implications of all models for equity index options. The implied volatilities are expressed as a function of moneyness. The empirical implied volatilities decline in moneyness, a pattern known as the implied volatility skew. The DA implied volatilities for the 1-month maturity are flat and approximately equal to the realized stock market volatility. One apparent candidate to generate a steep volatility skew is high risk aversion. Although raising risk aversion in Epstein–Zin preferences improves the model’s performance, this cannot fully account for the level of implied volatilities. In contrast, the GDA framework can fit the option prices much better. Figure 8 additionally presents implied volatilities for ATM and 0.90 OTM options. In the data, ATM (OTM) volatilities slightly increase (decrease) over the horizon. Neither DA nor EZ specification can match the level of the empirical curves. In contrast, GDA can explain overall patterns and magnitudes of the empirical implied volatilities.

FIGURE 8
Implied Volatilities

Graph A of Figure 8 plots the 1-month implied volatility curve as a function of moneyness for the data and the three models: GDA, DA, and EZ. Graphs B and C plot the empirical and model-based implied volatility curves for ATM and OTM options as functions of the time to maturity (in months). The empirical statistics are for the U.S. data from Jan. 1996 to Dec. 2016. The model-based curves are calculated for option prices using the annualized model-implied interest rate and dividend yield. For each model, I simulate 10,000 economies at a monthly frequency with a sample size equal to its empirical counterpart and report the medians of sample statistics. The model-implied results are based on the simulations without consumption disasters, consistent with the historical data.



V. Conclusion

I build an equilibrium model with GDA preferences and rare events in consumption growth. I show that the combination of the investor's tail aversion and fluctuating economic uncertainty due to learning about a hidden depression state explains a wide variety of asset pricing phenomena. Most notably, the model rationalizes the variance term structure, a new stylized fact of the variance swap data. In particular, the model predicts large and negative Sharpe ratios on 1-month variance forwards and produces a slightly positive term structure for maturities longer than 2 months. Furthermore, the model accounts for the large variance and skew risk premiums, and generates a realistic volatility surface implied by index options, while simultaneously matching the salient features of equity returns and the risk-free rate. I show that the success of the model is attributable to GDA by comparing GDA preferences to nested utilities: disappointment aversion and Epstein–Zin preferences. Although three specifications can reasonably match equity moments, only GDA preferences can explain the variance term structure, moment risk premiums, and option prices.

There are several interesting avenues for future research. First, my article highlights the importance of the specific values of the disappointment threshold and disappointment aversion. Although Delikouras (2017) provides the empirical estimate of a disappointment aversion parameter in Gul (1991), joint estimation of the parameters in Routledge and Zin (2010) has not been addressed yet. Second, it is fruitful to explore the implications of the richer model for the term structure of dividend strips and interest rates. For instance, the extension with post-depression recoveries (Hasler and Máfe (2016)) has the potential to jointly explain the term structures of interest rates, equity and variance risk. Third, GDA is likely to have additional asset pricing implications for risk premia with a multidimensional-learning problem (Johannes et al. (2016)) or rational parameter learning (Collin-Dufresne et al. (2016)). Finally, it is interesting to study GDA preferences with other behavioral biases (Brandt, Zeng, and Zhang (2004)).

Supplementary Material

To view supplementary material for this article, please visit <http://doi.org/10.1017/S0022109023000364>.

References

- Ai, H.; M. M. Croce; A. M. Diercks; and K. Li. "News Shocks and the Production-Based Term Structure of Equity Returns." *Review of Financial Studies*, 31 (2018), 2423–2467.
- Andrei, D.; M. Hasler; and A. Jeanneret. "Asset Pricing with Persistence Risk." *Review of Financial Studies*, 32 (2019), 2809–2849.
- At-Sahalia, Y.; M. Karaman; and L. Mancini. "The Term Structure of Equity and Variance Risk Premia." *Journal of Econometrics*, 219 (2020), 204–230.
- Augustin, P., and R. Tédongap. "Real Economic Shocks and Sovereign Credit Risk." *Journal of Financial and Quantitative Analysis*, 51 (2016), 541–587.
- Augustin, P., and R. Tédongap. "Disappointment Aversion, Term Structure, and Predictability Puzzles in Bond Markets." *Management Science*, 67 (2021), 6266–6293.

- Babiak, M., and R. Kozhan. "Growth Uncertainty, Rational Learning, and Option Prices." Working Paper. Lancaster University and Warwick Business School (2020).
- Babiak, M., and R. Kozhan. "Parameter Learning in Production Economies." Working Paper, Lancaster University and Warwick Business School (2021).
- Backus, D.; M. Chernov; and I. Martin. "Disasters Implied by Equity Index Options." *Journal of Finance*, 66 (2011), 1969–2012.
- Bakshi, G.; N. Kapadia; and D. Madan. "Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options." *Review of Financial Studies*, 16 (2003), 101–143.
- Bansal, R., and A. Yaron. "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles." *Journal of Finance*, 59 (2004), 1481–1509.
- Barro, R. J. "Rare Disasters and Asset Markets in the Twentieth Century." *Quarterly Journal of Economics*, 121 (2006), 823–866.
- Belo, F.; P. Collin-Dufresne; and R. S. Goldstein. "Dividend Dynamics and the Term Structure of Dividend Strips." *Journal of Finance*, 70 (2015), 1115–1160.
- Benzoni, L.; P. Collin-Dufresne; and R. S. Goldstein. "Explaining Asset Pricing Puzzles Associated with the 1987 Market Crash." *Journal of Financial Economics*, 101 (2011), 552–573.
- Berger, D.; I. Dew-Becker; and S. Giglio. "Uncertainty Shocks as Second-Moment News Shocks." *Review of Economic Studies*, 87 (2020), 40–76.
- Berrada, T.; J. Detemple; and M. Rindisbacher. "Asset Pricing with Beliefs-Dependent Risk Aversion and Learning." *Journal of Financial Economics*, 128 (2018), 504–534.
- Bollerslev, T.; G. Tauchen; and H. Zhou. "Expected Stock Returns and Variance Risk Premia." *Review of Financial Studies*, 22 (2009), 4463–4492.
- Bonomo, M.; R. Garcia; N. Meddahi; and R. Tédongap. "Generalized Disappointment Aversion, Long-Run Volatility Risk, and Asset Prices." *Review of Financial Studies*, 24 (2011), 82–122.
- Bonomo, M.; R. Garcia; N. Meddahi; and R. Tédongap. "The Long and the Short of the Risk-Return Trade-Off." *Journal of Econometrics*, 187 (2015), 580–592.
- Brandt, M. W.; Q. Zeng; and L. Zhang. "Equilibrium Stock Return Dynamics Under Alternative Rules of Learning About Hidden States." *Journal of Economic Dynamics and Control*, 28 (2004), 1925–1954.
- Cecchetti, S. G.; P. S. Lam; and N. C. Mark. "Mean Reversion in Equilibrium Asset Prices." *American Economic Review*, 80 (1990), 398–418.
- Choi, H.; P. Mueller; and A. Vedolin. "Bond Variance Risk Premiums." *Review of Finance*, 21 (2017), 987–1022.
- Choi, S.; R. Fisman; D. Gale; and S. Kariv. "Consistency and Heterogeneity of Individual Behavior Under Uncertainty." *American Economic Review*, 97 (2007), 1921–1938.
- Cogley, T., and T. J. Sargent. "The Market Price of Risk and the Equity Premium: A Legacy of the Great Depression?" *Journal of Monetary Economics*, 55(3) (2008), 454–476.
- Collin-Dufresne, P.; M. Johannes; and L. A. Lochstoer. "Parameter Learning in General Equilibrium: The Asset Pricing Implications." *American Economic Review*, 106 (2016), 664–698.
- Croce, M. M.; M. Lettau; and S. C. Ludvigson. "Investor Information, Long-Run Risk, and the Term Structure of Equity." *Review of Financial Studies*, 28 (2014), 706–742.
- Dahlquist, M.; A. Farago; and R. Tédongap. "Asymmetries and Portfolio Choice." *Review of Financial Studies*, 30 (2016), 667–702.
- David, A. "Fluctuating Confidence in Stock Markets: Implications for Returns and Volatilities." *Journal of Financial and Quantitative Analysis*, 32 (1997), 427–462.
- Delikouras, S. "Where's the Kink? Disappointment Events in Consumption Growth and Equilibrium Asset Prices." *Review of Financial Studies*, 30 (2017), 2851–2889.
- Delikouras, S., and A. Kostakis. "A Single-Factor Consumption-Based Asset Pricing Model." *Journal of Financial and Quantitative Analysis*, 54 (2019), 789–827.
- Dew-Becker, I., and S. Giglio. "Cross-Sectional Uncertainty and the Business Cycle: Evidence from 40 Years of Options Data." NBER Working Paper No 27864 (2020).
- Dew-Becker, I.; S. Giglio; and B. Kelly. "Hedging Macroeconomic and Financial Uncertainty and Volatility." *Journal of Financial Economics*, 142 (2021), 23–45.
- Dew-Becker, I.; S. Giglio; A. Le; and M. Rodriguez. "The Price of Variance Risk." *Journal of Financial Economics*, 123 (2017), 225–250.
- Drechsler, I. "Uncertainty, Time-Varying Fear, and Asset Prices." *Journal of Finance*, 68 (2013), 1837–1883.
- Drechsler, I., and A. Yaron. "What's Vol Got to Do With It." *Review of Financial Studies*, 24 (2011), 1–45.
- Du, D. "General Equilibrium Pricing of Options with habit formation and event Risks." *Journal of Financial Economics*, 99 (2011), 400–426.

- Epstein, L. G., and S. E. Zin. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica*, 57 (1989), 937–969.
- Eraker, B., and I. Shaliastovich. "An Equilibrium Guide to Designing Affine Pricing Models." *Mathematical Finance*, 18 (2008), 519–543.
- Farago, A., and R. Tédongap. "Downside Risks and the Cross-Section of Asset Returns." *Journal of Financial Economics*, 129 (2018), 69–86.
- Favilukis, J., and X. Lin. "Wage Rigidity: A Quantitative Solution to Several Asset Pricing Puzzles." *Review of Financial Studies*, 29 (2015), 148–192.
- Gabaix, X. "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance." *Quarterly Journal of Economics*, 127 (2012), 645–700.
- Gillman, M.; M. Kejak; and M. Pakos. "Learning About Rare Disasters: Implications for Consumption and Asset Prices." *Review of Finance*, 19 (2015), 1053–1104.
- Gul, F. "A Theory of Disappointment Aversion." *Econometrica*, 59 (1991), 667–686.
- Hansen, L. P. "Beliefs, Doubts and Learning: The Valuation of Macroeconomic Risk." *American Economic Review*, 97 (2007), 1–30.
- Hasler, M.; M. Khapko; and R. Márfe. "Rational Learning and Term Structures." Working Paper, University of Toronto and Collegio Carlo Alberto (2019).
- Hasler, M., and R. Márfe. "Disaster Recovery and the Term Structure of Dividend Strips." *Journal of Financial Economics*, 122 (2016), 116–134.
- Johannes, M.; L. A. Lochstoer; and Y. Mou. "Learning About Consumption Dynamics." *Journal of Finance*, 71 (2016), 551–600.
- Ju, N., and J. Miao. "Ambiguity, Learning, and Asset Returns." *Econometrica*, 80 (2012), 559–591.
- Judd, K. L. "Projection Methods for Solving Aggregate Growth Models." *Journal of Economic Theory*, 58 (1992), 410–452.
- Kozhan, R.; A. Neuberger; and P. Schneider. "The Skew Risk Premium in the Equity Index Market." *Review of Financial Studies*, 26 (2013), 2174–2203.
- Liu, H., and J. Miao. "Growth Uncertainty, Generalized Disappointment Aversion and Production-Based Asset Pricing." *Journal of Monetary Economics*, 69 (2014), 70–89.
- Liu, J.; J. Pan; and T. Wang. "An Equilibrium Model of Rare-Event Premia and its Implication for Option Smirks." *Review of Financial Studies*, 18 (2005), 131–164.
- Londono, J. M., and H. Zhou. "Variance Risk Premiums and the Forward Premium Puzzle." *Journal of Financial Economics*, 124 (2017), 415–440.
- Lorenz, F.; K. Schmedders; and M. Schumacher. "Nonlinear Dynamics in Conditional Volatility." In *Proceedings of Paris December 2020 Finance Meeting EUROFIDAI-ESSEC*; available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3575458 (2020).
- Márfe, R. "Income Insurance and the Equilibrium Term Structure of Equity." *Journal of Finance*, 72 (2017), 2073–2130.
- Melino, A., and A. X. Yang. "State-Dependent Preferences Can Explain the Equity Premium Puzzle." *Review of Economic Dynamics*, 6 (2003), 806–830.
- Nakamura, E.; J. Steinsson; R. Barro; and J. Ursua. "Crises and Recoveries in an Empirical Model of Consumption Disasters." *American Economic Journal: Macroeconomics*, 5 (2013), 135–174.
- Pastor, L., and P. Veronesi. "Learning in Financial Markets." *Annual Review of Financial Economics*, 1 (2009), 361–381.
- Pohl, W.; K. Schmedders; and O. Wilms. "Higher Order Effects in Asset Pricing Models with Long-Run Risks." *Journal of Finance*, 73 (2018), 1061–1111.
- Rietz, T. A. "The Equity Risk Premium: A Solution." *Journal of Monetary Economics*, 22 (1988), 117–131.
- Routledge, B. R., and S. E. Zin. "Generalized Disappointment Aversion and Asset Prices." *Journal of Finance*, 65 (2010), 1303–1332.
- Schreindorfer, D. "Macroeconomic Tail Risks and Asset Prices." *Review of Financial Studies*, 33 (2020), 3541–3582.
- Seo, S. B., and J. A. Wachter. "Option Prices in a Model with Stochastic Disaster Risk." *Management Science*, 65 (2019), 3449–3469.
- Shaliastovich, I. "Learning, Confidence, and Option Prices." *Journal of Econometrics*, 187 (2015), 18–42.
- van Binsbergen, J.; M. Brandt; and R. Kojien. "On the Timing and Pricing of Dividends." *American Economic Review*, 102 (2012), 1596–1618.
- van Binsbergen, J.; W. Hueskes; R. Kojien; and E. Vrugt. "Equity Yields." *Journal of Financial Economics*, 110 (2013), 503–519.
- van Binsbergen, J. H., and R. S. Kojien. "The Term Structure of Returns: Facts and Theory." *Journal of Financial Economics*, 124 (2017), 1–21.

- Veronesi, P. "Stock Market Overreaction to Bad News in Good Times: A Rational Expectations Equilibrium Model." *Review of Financial Studies*, 12 (1999), 975–1007.
- Veronesi, P. "How Does Information Quality Affect Stock Returns." *Journal of Finance*, 155 (2000), 807–837.
- Weitzman, M. "Subjective Expectations and Asset-Return Puzzles." *American Economic Review*, 97 (2007), 1102–1130.
- Zhou, G., and Y. Zhu. "Macroeconomic Volatilities and Long-Run Risks of Asset Prices." *Management Science*, 61 (2014), 413–430.