On Sources of Risk Premia in Representative Agent Models*

Tyler Beason David Schreindorfer

Arizona State University

February 5, 2020

Abstract

We use options and return data to decompose unconditional risk premia into different parts of the return state space. In the data, the entire equity premium is attributable to monthly returns *below* -11.3%, but returns in the extreme left tail matter very little. In contrast, leading asset pricing models based on habits, long-run risks, and rare disasters attribute the premium almost exclusively to returns *above* -11.3%, or to the extreme left tail. We find that model extensions with a larger quantity of tail risk cannot account for the data, while models with a higher price of tail risk can.

KEYWORDS: Equity premium puzzle, equity index options, Arrow-Debreu securities, price of risk, sources of risk premia, model diagnostic, representative agent models, rare disasters, long-run risks, external habits

*Contact: tyler.beason@asu.edu and david.schreindorfer@asu.edu. We appreciate helpful comments and suggestions from David Bates, Sreedhar Bharath, Michael Barnett, Oliver Boguth, Philip Bond, Christopher Hrdlicka, Seth Pruitt, Sunil Wahal; seminar participants at Arizona State University, Carnegie Mellon University, the Federal Reserve Board, the University of Iowa, and the University of Washington.

1 Introduction

Most economists would agree that macroeconomic risk is an important source of stock market risk premia. Indeed, equilibrium asset pricing theories typically assume that fundamental shocks represent *the only source* of priced risk, as in Lucas (1978). There is far less agreement, however, about the types of macroeconomic shocks reflected in asset prices or investors' risk attitudes towards them. An important reason is that theories based on vastly different fundamental risks, such as Campbell and Cochrane (1999), Bansal and Yaron (2004), Rietz (1988), and Barro (2006), have similar implications for many key data features, including the average equity premium.

In this paper, we propose a data-based metric for discriminating between asset pricing theories based on *how* they generate the equity premium. Our metric builds on the concept of Arrow-Debreu securities – claims to a unit payoff in a particular state of nature. We show that the absence of arbitrage opportunities allows us to write the unconditional equity premium as

$$\mathbb{E}[R_{t+1} - R_t^f] = \int_{-1}^{\infty} R\,[f(R) - f^*(R)]\,dR, \tag{1}$$

where $f(R)$ and $f^*(R) = \mathbb{E}[f_t^*(R)]$ are, respectively, the average payoff and average forward price of an Arrow-Debreu security that pays \$1 if the realized stock market return equals R. The average payoff can equivalently be interpreted as an unconditional return distribution. The forward price, $f_t^*(R)$, can equivalently be interpreted as a preference-adjusted, or "risk-neutralized", return distribution under which the expected stock market return equals the risk-free rate. Equation (1) represents a powerful tool for understanding how stock prices are formed. Specifically, it allows us to assess how much individual states contribute to the overall equity premium by evaluating the integral separately for different regions.

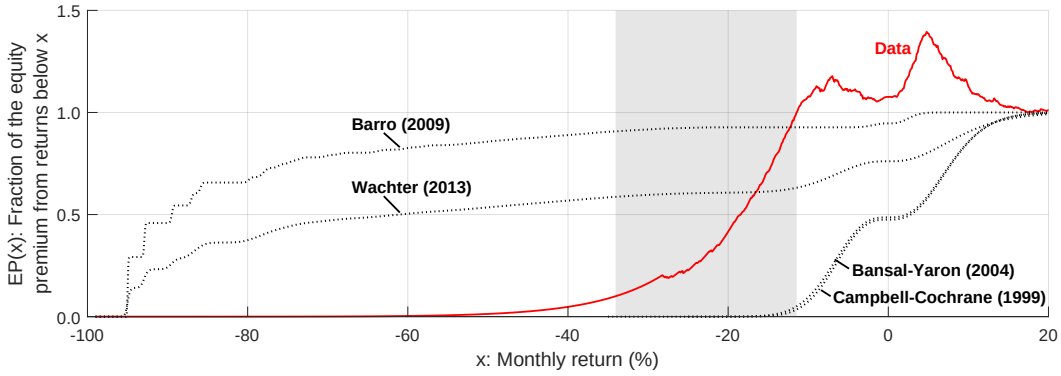

Figure I: Sources of the equity premium. The integral in Equation (1) is shown as a function of the upper limit of integration, and normalized by the total equity premium: $EP(x) \equiv \frac{1}{\mathbb{E}[R_{t+1}-R_t^f]} \int_{-1}^{x} R\left[f(R) - f^*(R)\right] dR$. The sample spans 1990-2018. The shaded area marks monthly returns between -34% and -11.3%, a set of states that accounts for 9/10 of the equity premium in the data.

We implement the decomposition empirically for the S&P 500, a broad US stock market index, by estimating f from returns and f^* from option prices. Using these estimates, the solid line in Figure I shows the fraction of the equity premium associated with returns below different thresholds, a function we coin "$EP(x)$". By construction, increasing segments of the $EP(x)$ curve reflect states that contribute positively to the equity premium, whereas decreasing segments reflect states that contribute negatively. Strikingly, the figure shows that the entire equity premium is associated with monthly returns below -11.3%, while 9/10 of it stems from the gray-shaded region between -34% and -11.3%. Such returns occur roughly at a business cycle frequency, and they have historically coincided with events like Iraq's invasion of Kuwait in 1990, the collapse of Long Term Capital Management in 1998, the September 11 attacks in 2001, and the bankruptcy of Lehman Brothers in 2008. Investors are therefore predominantly compensated for shocks that coincide with infrequent, large, but not extremely large negative stock market returns.

We show that the shape of $EP(x)$ to the right of the gray-shaded region, including the fact that the curve rises above one, is consistent with a pricing kernel that is U-shaped when projected onto returns. Because this non-monotonicity has been extensively documented in prior work and mechanisms to explain it exist, we do not make it a focal point of our paper.[1]

The remaining lines in Figure I show $EP(x)$ for different asset pricing theories. The habit model of Campbell and Cochrane (1999) and the long-run risks model of Bansal and Yaron (2004) build on very different economic mechanisms, but nevertheless attribute the equity premium to nearly identical return states. In both models, positive returns account for about half of the equity premium, while returns below -11.3% account for very little. We show that these features are a result of lognormality. The rare disaster model of Barro (2006, 2009) attributes 91/100 of the equity premium to states on the left of the gray-shaded region, and more than half to returns below -89%. By extending the Barro model with time variation in the probability of disasters and a recursive utility agent, Wachter (2013) combines the long-run risks and disaster mechanisms. We find that her model attributes about 2/3 of the equity premium to the far left tail and 1/3 to states on the right of the gray-shaded region. This evidence suggests that sources of stock market risk premia in popular asset pricing models differ substantially from those in the data.

To understand the origin of these discrepancies, consider the Arrow-Debreu security associated with a large negative return, $\tilde{R}$, in the gray-shaded region of Figure I. Because this security provides a valuable hedge to investors, its average forward price exceeds its average payoff, $f^*(\tilde{R}) > f(\tilde{R})$, and the state contributes positively to the equity premium. This effect is present in all of the aforementioned models, but its magnitude is considerably smaller than in the data. Equation (1) implies

[1] Cuesdeanu and Jackwerth (2018) provide a recent survey of this literature.

that the reason must either lie in $\tilde{R}$ being less likely than in the data, or in it being associated with lower levels of marginal utility – see Section 2 for details.

Recent studies have augmented the aforementioned models with left-skewed shocks in fundamentals to align them with option prices and higher moments of stock market returns. Naturally, doing so increases the probability of large negative returns. We find that returns in the gray-shaded region of Figure I account for approximately 1/3 of the equity premium in the habit model of Bekaert and Engstrom (2017), the long-run risks model of Drechsler and Yaron (2011), and the "frequent jump model" of Backus et al. (2011). This number is higher than in the original models, but differences between it and the empirical value of 9/10 remain economically large and statistically significant for all three extensions. We also find that the augmented models inflate the probability of the associated states by factors of 1.5 to 2.5 relative to the data. The scope for rationalizing our empirical evidence based on a larger *quantity of tail risk* alone therefore seems limited.

Backus et al. (2011) and Martin (2017) show that many representative agent models are inconsistent with option prices, in that they either imply too little or too much skewness in the option-implied return distribution. Our evidence is conceptually distinct from this finding because risk premia reflect *differences* between the option-implied and historical return distributions. Indeed, our analysis of the Backus et al. (2011) model shows that consistency with option prices does not imply consistency with $EP(x)$.

We argue that standard representative agent endowment economies provide only two opportunities to increase the *price of tail risk*, i.e. the average value of marginal utility coinciding with states like $\tilde{R}$. The first option is to alter the correlation structure between shocks to consumption and dividends. For example, if large drops in the stock market are likely to coincide with declines in aggregate economic activity (and hence elevated marginal utility), while small returns are more or less indepen-

dent of the real economy, the price of risk for tail events is relatively higher. The second option is to assume a utility function that places a greater emphasis on tail events. We provide examples of both channels in simple IID settings, and show that they can account for key features of the empirical $EP(x)$ curve. The models we consider are kept highly stylized to illustrate their mechanisms in the clearest possible way, i.e. we do not view them as fully specified asset pricing theories. Nevertheless, we believe that they illustrate the necessary ingredients for aligning representative agent models more closely with sources of the equity premium.

Our findings complement those of van Binsbergen et al. (2012), who rely on options data to show that near term dividends earn larger risk premia than far term dividends. Their evidence speaks to the pricing of shocks with different persistence levels and is therefore primarily informative about investors' time preferences. Our evidence speaks to the pricing of shocks with different magnitudes and is therefore primarily informative about investors' risk preferences. A complete explanation of the equity premium must be consistent with the pricing of all shocks to aggregate dividends: transient, persistent, small, and large.

We also relate closely to Bollerslev and Todorov (2011), who quantify the importance of "jumps" for the *conditional* equity premium in the data, and conclude that "*any satisfactory equilibrium-based asset pricing model must be able to generate large and time-varying compensations for the possibility of rare disasters*". Methodologically, our approach differs from theirs because we study the *unconditional* equity premium. As a result, we do not require an extreme value theory to approximate the probability of large low-frequency (e.g. monthly) returns based on the probability of large high-frequency (e.g. 5 minute) returns. Economically, we reach the opposite conclusion regarding the importance of rare disasters. In particular, our evidence shows that the disaster models of Barro (2009) and Wachter (2013) attribute the equity premium primarily to the extreme left tail, while the data attributes it almost

exclusively to an intermediate left tail region.

Welch (2018) shows that combining the market with a protective put that is 15% out-of-the-money lowers average returns by only 2% p.a. Based on this evidence, he argues that rare disasters account for at most 2% (in absolute terms) of the equity premium. Seo and Wachter (2018) show that this reasoning is incorrect, however, because the disaster model of Wachter (2013) replicates Welch's empirical evidence. Our paper is therefore the first to show that popular disaster models fail to account for sources of the equity premium.

More broadly, a large literature in empirical option pricing shows that investors are willing to pay substantial premia to hedge against big stock market corrections, a phenomenon Rubinstein (1994) coined "crash-o-phobia". Coval and Shumway (2001), Carr and Wu (2009), and Broadie et al. (2009) document that option-based hedging strategies generate Sharpe ratios that are negative and 2-5 times larger in magnitude than that of the market. More recently, Dew-Becker et al. (2017) show that leading representative agent models are inconsistent with the high Sharpe ratios of short maturity option portfolios. We build on this prior evidence by quantifying the importance of tail risk for the equity premium.

2 Decomposing the equity premium

We begin by deriving $EP(x)$ and relating it to the pricing kernel. Let $R_{t+1}^{cum} = \frac{S_{t+1}+D_{t+1}}{S_t}$ be the cum-dividend return on the market, $f_t(R)$ its conditional probability density function (PDF), and $f_t^*(R)$ the forward price of an Arrow-Debreu security that pays \$1 if $R_{t+1}^{cum} = R$ and \$0 otherwise. Note that the conditional PDF can be interpreted as the expected payoff of the Arrow-Debreu security for state R.

2.1 From Arrow-Debreu to the equity premium

A \$1 investment in the market has a payoff of R_{t+1}^{cum}, and can therefore be replicated with a portfolio that contains R units of the Arrow-Debreu security for state R, for all $R \in [0, \infty]$. Because the forward price of \$1 equals R_t^f and the forward price of the aforementioned portfolio equals $\int_0^\infty R f_t^*(R) dR$, the absence of arbitrage opportunities implies that[2]

$$\int_0^\infty R f_t^*(R) dR = R_t^f . \tag{2}$$

Using this equation, we can write the conditional equity premium as

$$\mathbb{E}_t[R_{t+1}^{cum}] - R_t^f = \int_0^\infty R \left[f_t(R) - f_t^*(R) \right] dR \tag{3}$$

and interpret it as the difference between the expected payoff and the forward price of a portfolio of Arrow-Debreu securities.

Two issues hinder our ability to estimate the integrant in Equation (3) empirically. First, option prices only allow us to identify Arrow-Debreu state prices for ex-dividend, rather than cum-dividend, returns. Second, estimates of the conditional return distribution require strong statistical assumptions and necessarily suffer from a mismatch between investors' and the econometrician's information set.

We address issue one by assuming that time-$(t+1)$ dividends are in investors' time-t information set. This assumption is reasonable for short return horizons such as one month, which are the focus of our empirical analysis, because S&P 500 companies announce dividends on average three weeks in advance of the ex-dividend

[2] "No arbitrage" also ensures that $f_t^*(R) \geq 0$ and $\int f_t^*(R) dR = 1$, so that $f_t^*(R)$ can equivalently be interpreted as a PDF. Equation 2 is therefore often written as $\mathbb{E}_t^*[R_{t+1}^{cum}] = R_t^f$ and interpreted as stating that the expected return equals the risk-free rate under the "risk-neutral" return distribution.

date (Schulz 2016, Table I).[3] It implies that the expected dividend yield and the dividend yield's forward price are identical, so that their difference drops out of Equation (3). We can therefore evaluate the integral based on expected payoffs and state prices for ex-dividend returns. In what follows, we work with net, rather than gross returns, because doing so leads to a more natural interpretation of our equity premium decomposition. This change amounts to subtracting the constant 1 from R, which also drops out of Equation (3) because $\int f = \int f^* = 1$. Henceforth, we therefore interpret f_t and f_t^* as the expected payoff and forward price of Arrow-Debreu securities that are written on the *net ex-dividend return*. In an abuse of notation, we continue to rely on the same symbols.

To address issue two, we take the unconditional expectation of Equation (3) and note that $\mathbb{E}[f_t(R)]$ equals the unconditional PDF, $f(R)$. Doing so yields Equation (1) in the introduction. To decompose the equity premium, we define

$$EP(x) \equiv \frac{\int_{-1}^{x} R\,[f(R) - f^*(R)]\,dR}{\int_{-1}^{\infty} R\,[f(R) - f^*(R)]\,dR}, \tag{4}$$

which measures the fraction of the unconditional equity premium that is associated with returns less than x. Note that the total equity premium in the denominator of (4) is a normalization constant that does not depend on x. The normalization ensures that, like a CDF, $EP(x)$ approaches zero for return thresholds in the far left tail and approaches one for return thresholds in the far right tail. Unlike a CDF, however, theory does not restrict $EP(x)$ to be monotonically increasing.

[3] van Binsbergen et al. (2012) show that claims to "near-term" dividends earn larger risk premia than the market itself. Because this finding is based on claims to market dividends over the coming 1.6 years (on average), however, it is not inconsistent with the absence of risk premia on very short-term dividends.

2.2 Relating $EP(x)$ to the pricing kernel

To interpret $EP(x)$ through the lens of equilibrium asset pricing models, it is useful to relate it to the pricing kernel, M. The absence of arbitrage opportunities implies

$$\frac{f_t^*(R_{t+1})}{f_t(R_{t+1})} = \frac{\mathbb{E}_t[M_{t+1}|R_{t+1}]}{\mathbb{E}_t[M_{t+1}]}. \tag{5}$$

The "projected pricing kernel", $\mathbb{E}_t[M_{t+1}|R_{t+1}]$, reveals how marginal utility varies with returns and it has the same pricing implications as M for the market and for any derivative that is written on the market.[4] Using (5), the conditional equity premium in (3) can be expressed as

$$\mathbb{E}_t[R_{t+1}] - R_t^f = \int_{-1}^{\infty} R f_t(R) \left(1 - \frac{\mathbb{E}_t[M_{t+1}|R]}{\mathbb{E}_t[M_{t+1}]}\right) dR. \tag{6}$$

Conditionally, returns are therefore important for risk premia if they are (i) large in magnitude (large absolute R), (ii) conditionally likely (high f_t), or (iii) negative and associated with high marginal utility ($\mathbb{E}_t[M_{t+1}|R_{t+1}] > E_t[M_{t+1}]$), or positive and associated with low marginal utility ($\mathbb{E}_t[M_{t+1}|R_{t+1}] < E_t[M_{t+1}]$). Taking the unconditional expectation of (6) and using the fact that $\mathbb{E}[XY] = \mathbb{E}[X]\mathbb{E}[Y] + cov(X,Y)$ shows that the unconditional equity premium equals

$$\begin{aligned} &\mathbb{E}[R_{t+1} - R_t^f] \\ &= \int_{-1}^{\infty} R \times \left(f(R) \left(1 - \mathbb{E}\left[\frac{\mathbb{E}_t[M_{t+1}|R]}{\mathbb{E}_t[M_{t+1}]}\right]\right) - cov\left[f_t(R), \frac{\mathbb{E}_t[M_{t+1}|R]}{\mathbb{E}_t[M_{t+1}]}\right]\right) dR. \end{aligned} \tag{7}$$

Unconditionally, returns are important for risk premia if they are (i) large in magnitude, (ii) unconditionally likely, (iii) negative and associated with high average marginal utility or positive and associated with low average marginal utility, or (iv) negative and have a conditional probability, $f_t(R)$, that co-varies positively with their conditional marginal utility, $\mathbb{E}_t[M_{t+1}|R]/E_t[M_{t+1}]$, or if they are positive

[4] Specifically, the law of iterated expectations ensures that $\mathbb{E}_t[M_{t+1}g(R_{t+1})] = \mathbb{E}_t[\mathbb{E}_t[M_{t+1}|R_{t+1}]g(R_{t+1})]$ for any payoff function $g(\cdot)$.

and have a conditional probability that co-varies negatively with their conditional marginal utility. The covariance term equals zero in IID environments. However, we find that it is also quantitatively negligible relative to the first term in (7) in the non-IID models we consider (Campbell-Cochrane 1999, Bansal-Yaron 2004, Drechsler-Yaron 2011, Wachter 2013, and Bekaert-Engstrom 2017). These results are detailed in the online appendix.

Equation (7) helps us understand two features of the empirical $EP(x)$ curve in Figure I. First, the curve is falling for small negative returns and rising for small positive returns. Ignoring the covariance term in (7), this implies that returns in this region are associated with below average marginal utility, $\mathbb{E}\left[\frac{\mathbb{E}_t[M_{t+1}|R]}{\mathbb{E}_t[M_{t+1}]}\right] < 1$. Second, the curve is falling for large positive returns. Because $EP(\infty) = 1$, this feature is equivalent to the observation that the $EP(x)$ curve rises above one. Again ignoring the covariance term, it implies that large positive returns are associated with above average marginal utility, $\mathbb{E}\left[\frac{\mathbb{E}_t[M_{t+1}|R]}{\mathbb{E}_t[M_{t+1}]}\right] > 1$. Taken together, these observations imply that the $EP(x)$ curve is consistent with a U-shaped pricing kernel.

3 Estimation

We discuss data sources and summarize our approach for estimating average payoffs and forward prices of Arrow-Debreu securities. The appendix contains additional details about data filters, our approach for estimating f^*, and robustness tests.

3.1 Payoffs

The payoff of an Arrow-Debreu security for state R is given by the indicator function $\mathbf{1}\{R_{t+1} = R\}$. Its conditional expectation therefore equals the conditional PDF, $\mathbb{E}_t[\mathbf{1}\{R_{t+1} = R\}] = f_t(R)$. We estimate the average payoffs, $f(R) = \mathbb{E}[f_t(R)]$, with the unconditional empirical PDF, which attaches a probability of $1/T$ to each return

observation. The sample is daily, spans 1990-2018, and consists of overlapping 30 calendar day ex-dividend returns on the S&P 500 index. Our estimate of $f(R)$ implies that monthly returns have an annualized volatility of 15.02%, skewness of -0.69, kurtosis of 7.22, and a probability of 1.08% of falling below -11.3%. Returns below -11.3% have therefore occurred every $\frac{1}{12\times 0.0108} = 7.7$ years on average, whereas returns below -34% have not occurred in our sample.

3.2 Forward prices

We rely on a well-known result by Breeden and Litzenberger (1978) to recover f_t^* from options prices. Specifically, the absence of arbitrage opportunities implies that the price of a European put option with strike K and moneyness $X = K/S_t$ equals

$$P_t(X) = \frac{S_t}{R_t^f} \int_{-1}^{X} (X - R) f_t^*(R) dR. \tag{8}$$

Differentiating this expression with respect to moneyness shows that the forward price of a portfolio of Arrow-Debreu securities for states below R (the conditional risk-neutral CDF), which we denote by $F_t^*(R) = \int_{-1}^{R} f_t^*(X) dX$, equals

$$F_t^*(R) = S_t \times R_t^f \times \left. \frac{\partial P_t(X)}{\partial X} \right|_{X=R}. \tag{9}$$

Differentiating once again with respect to X shows that forward state prices (the conditional risk-neutral PDF) are given by

$$f_t^*(R) = S_t \times R_t^f \times \left. \frac{\partial^2 P_t(X)}{\partial X^2} \right|_{X=R}. \tag{10}$$

Because $P_t(X)$, $F_t^*(R)$, and $f_t^*(R)$ can all be interpreted as prices, the absence of arbitrage opportunities dictates that put prices are (i) non-negative, (ii) monotonically increasing in moneyness, and (iii) convex in moneyness. When estimating forward state prices based on Equation (10) in the data, we impose these theoretical

constraints to discipline the shape of f_t^* at moneyness levels and maturities where few option price quotes are available.

Option price quotes for the period from January 5, 1990 to December 31, 2018 are obtained from the Chicago Board Options Exchange (CBOE). The dataset contains end-of-day information of all European exercise style options written on the S&P 500 index (underlying SPX). We apply standard filters to remove observations with low liquidity and obvious data errors and measure option prices by the average of bid and ask quotes.

In a perfect world, liquid option contracts would be available for a continuum of moneyness-maturity pairs over $[-1, \infty] \times [0, \infty]$, and $f_t^*(R)$ would be recoverable based on Equation (10) alone. In reality, option price quotes are somewhat noisy, only available for a discrete set of moneyness-maturity pairs, and contracts are not available for very long maturities or extreme moneyness levels. It is therefore necessary to smooth, interpolate, and extrapolate observed price quotes to obtain sensible estimates of f_t^*. Doing so requires some structure. In a nutshell, our approach fits a flexible, bivariate function in moneyness and maturity to observed option prices subject to theoretical no arbitrage constraints, evaluates the function at a maturity of 30 calendar days and a fine grid or moneyness, and computes state prices based on Equation (10) via finite differences. The appendix provides details. Relative to estimation methods that fit option prices separately for different maturities, our approach has the important advantage that it allows us to estimate 30-day state prices on every day of the sample, even those without 30-day option price quotes. Because 30-day quotes are only observed on one day per month for most of the sample, this feature makes our estimate of average state prices considerably more precise.

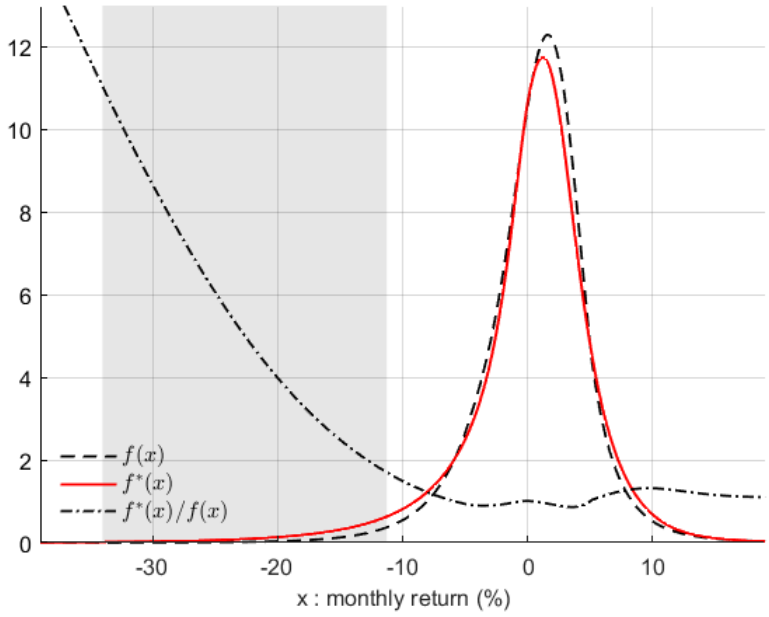

Figure II: Average Arrow-Debreu payoffs and forward prices in 1990-2018 data.

The estimated $f^*(R)$ is shown in Figure II, together with a smooth estimate of $f(R)$.[5] Under f^*, monthly returns have an annualized volatility of 19.1%, skewness of -1.56, kurtosis of 10.72, and a 3.50% probability of falling between -34% and -11.3% (one event every 2.4 years). This probability is 3.24 times higher than the corresponding historical probability. The dash-dotted line in Figure II shows same probability ratio for the entire return state space. Clearly, returns in the far left tail are associated with very high marginal utility for investors. In sections 4–6, we compare these probability ratios to their counterparts in asset pricing models.

[5]The empirical PDF consists of point masses and is therefore challenging to plot. For Figure II, we approximate f with a 10-th order polynomial between percentiles 10 and 90 and with a generalized extreme value (GEV) distribution in both extreme tails. The GEV distribution nests the tail shape of most parametric families of PDFs and has favorable asymptotic properties; see the extreme value theorem and discussion in Figlewski (2008). We use the empirical PDF (without smoothing) for all other results in the paper.

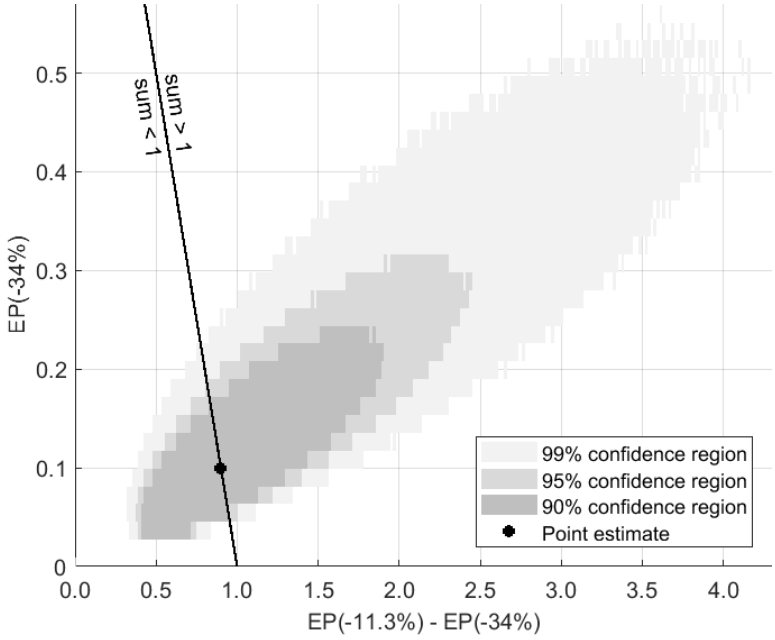

Figure III: Empirical sampling distribution of $EP(x)$. For a partition of the return state space, the figure shows the joint sampling distribution of each region's contribution to the equity premium (as a fraction of the total equity premium). The distribution is constructed with a block-bootstrap, using a block length of 21 trading days and 10 million bootstrap samples.

3.3 Statistical precision

Figure I shows $EP(x)$ for our empirical estimates of $f(R)$ and $f^*(R)$. To assess the amount of statistical uncertainty associated with key features of the curve, we use a block bootstrap to construct the joint sampling distribution of $EP(-34\%)$ and $EP(-11.3\%)-EP(-34\%)$. Figure III shows that it is very likely that most of the equity premium is attributable to return states between -34% and -11.3% (the gray-shaded region in Figure I) in population. In particular, $EP(-11.3\%) - EP(-34\%)$ exceeds one half in 97.8% of all bootstrap samples. Similarly, returns below -34% are very likely not the main source of stock market risk premia. $EP(-34\%)$ falls below one half in 98.2% of bootstrap samples. It is also worth noting that points to the

north-east of the solid line represent samples for which the estimated $EP(-11.3\%)$ exceeds one, so that returns above -11.3% contribute negatively to the equity premium. This occurs in 47.1% of samples.

One issue the bootstrap cannot account for is the potential underrepresentation of "disasters" in our 1990-2018 sample. If more disasters had occurred, however, $f(R)$ would have had *more* probability mass in the far left tail, the difference between f^* and f in that region would have been *smaller*, and such states would have accounted for *less* of the equity premium. If anything, the low importance of states below -34% therefore reflects the opposite of a "Peso problem".

4 $EP(x)$ in lognormal asset pricing models

Many asset pricing models are lognormal – they imply joint normality of log returns and the log pricing kernel – and yield both tractable solutions and simple intuition as a result. It is well-known, however, that log returns for most assets are imperfectly approximated by the normal distribution. We evaluate what lognormality implies about our state space decomposition of the equity premium in a simple, IID representative agent model. Next, we show that conditionally lognormal models with richer dynamics, such as Campbell and Cochrane (1999) and Bansal and Yaron (2004), have almost identical implications for $EP(x)$. The appendix discusses how $EP(x)$ is computed for all models under consideration.

4.1 Unconditional lognormality

We begin with a description of the IID model (henceforth "Model 1"). 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. Log consumption growth $\Delta c_{t+1} =$

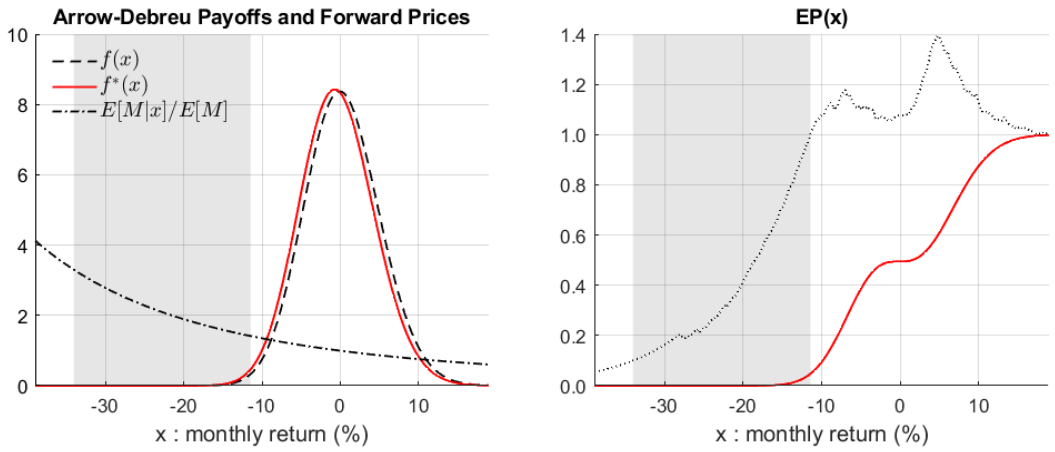

Figure IV: Unconditional lognormality (Model 1).

$\ln\left(\frac{C_{t+1}}{C_t}\right)$ and log dividend growth $\Delta d_{t+1} = \ln\left(\frac{D_{t+1}}{D_t}\right)$ are given by

$$\begin{aligned} \Delta c_{t+1} &= g + \sigma \varepsilon^c_{t+1} \\ \Delta d_{t+1} &= g + \varphi\sigma \varepsilon^d_{t+1}, \end{aligned} \tag{11}$$

where ε^c and ε^d are bivariate standard normal with correlation $\varrho = 1$ and IID over time. Section 6 considers cases with imperfect correlation. Our monthly calibration assumes mean growth rates of $g = 0.02/12$, consumption growth volatility of $\sigma = 0.03/\sqrt{12}$, and "leverage" of $\varphi = 5.5$. The representative investor has power utility preferences with a time discount factor of $\beta = 0.99^{1/12}$. We calibrate the coefficient of relative risk aversion, $\gamma = 15.84$, to generate an equity premium of 8% per year.

The IID assumption implies $\ln(1 + R_t) = \Delta d_t$, so that average Arrow-Debreu payoffs (f) are given by a lognormal PDF with parameters g and $\varphi\sigma$. The power utility function and lognormal endowment imply that Arrow-Debreu forward prices (f^*) equal a lognormal PDF with parameters $g - \gamma\varrho\varphi\sigma^2$ and $\varphi\sigma$ (the online appendix provides a derivation). The risk-neutral density of log returns therefore has the same variance and a lower mean than the historical density – it is "shifted left".

Figure IV shows $EP(x)$ for Model 1, along with its determinants f and f^*.

About half of the equity premium is attributable to positive and negative returns, respectively. The reason is that f is nearly symmetric with a mean around zero, while the pricing kernel overweights negative returns by about as much as it underweights positive returns. The risk-neutral density of simple returns, f^*, therefore resembles a "left-shifted" version of f, and $\int R(f - f^*)$ is about equal for positive and negative returns. The figure further shows that states in the gray-shaded region account for little of the equity premium. An obvious reason lies in the small quantity of tail risk under lognormality. Table II shows that such states have a probability of 0.005, less than half their empirical probability of 0.011. We quantify the price of tail risk with the ratio of the average risk-neutral probability of states in the gray-shaded region and their historical probability, $\int_{-34\%}^{-11.3\%} f^*(R)dR / \int_{-34\%}^{-11.3\%} f(R)dR$. The numerator equals the average forward price of an asset that pays \$1 if the realized stock market return falls into the gray-shaded region, and the denominator equals its average payoff. Table II shows that the lognormal model implies a price of tail risk of 1.468, compared to a value of 3.236 in the data. Because the IID assumption implies that $\mathbb{E}[M|R]/\mathbb{E}[M] = f^*/f$ unconditionally, we can also directly compare the pricing kernel in Figure IV to the empirical density ratio in Figure II. Clearly, the lognormal model implies that Arrow-Debreu forward prices in the far left tail are substantially lower than in the data. Hence, the lognormal model understates the importance of left tail events because the quantity and price of tail risk are both substantially lower than in the data.

4.2 Habits and long-run risks

The habit model of Campbell and Cochrane (1999) and the long-run risks model of Bansal and Yaron (2004) are approximately conditionally lognormal. Figure I shows that their implied $EP(x)$ curves look almost identical to that in the IID lognormal

Table I: Characteristics of important return states: $R \in [-34\%, -11.3\%]$

	EP fraction	$\int f(R)dR$	$\frac{\int f^*(R)dR}{\int f(R)dR}$
Data, 1990-2018	0.900	0.011	3.236
Lognormality			
Model 1	0.052^{***}	0.005^{**}	1.468^{*}
Campbell-Cochrane (1999)	0.038^{***}	0.004^{**}	1.391^{*}
Bansal-Yaron (2004)	0.055^{**}	0.006	1.311^{*}
Increased quantity of tail risk			
Model 2	0.271^{**}	0.011	1.755
Bekaert-Engstrom (2017)	0.369^{**}	0.028^{**}	1.511^{**}
Drechsler-Yaron (2011)	0.357^{**}	0.017	1.738
Backus-Chernov-Martin (2011)	0.304^{*}	0.021^{*}	1.378^{**}
Barro (2009)	0.020^{***}	0.001^{***}	1.676
Wachter (2013)	0.046^{***}	0.007	1.245
Increased price of tail risk			
Model 3	0.801	0.011	3.135
Model 4	0.967	0.005^{**}	9.935^{***}
Schreindorfer (2019)	0.720	0.008	4.685

"EP fraction": fraction of the equity premium associated with monthly returns between -34% and -11.3%; $\int f(R)dR$: probability of such returns; $\frac{\int f^*(R)dR}{\int f(R)dR}$: ratio of their average risk-neutral to actual probability. Small sample p-values are computed from 1 million simulated samples of length 336 months under each model. * $p < .1$, ** $p < .05$, *** $p < .01$.

model. Additionally, both models have very similar implications for the quantity and price of tail risk, both of which are lower than in the data (see Table II). As in the IID model, states in the gray-shaded tail region contribute about 5/100 to the equity premium. The empirical contribution of 9/10 has a small sample p-value below 5% in both models. Hence, time-variation in economic state variables has little effect on the models' pricing implications for shocks of different magnitudes.

5 Increasing the quantity of tail risk

To what extent does the unimportance of left tail events in lognormal models reflect their small probability? To what extent does it reflect a small price of risk? To answer these questions, we augment the lognormal model in Section 4.1 with left-skewed shocks in fundamentals. As in Backus et al. (2011), consumption and dividends are subject to both small shocks and "jumps",

$$\begin{aligned} \Delta c_{t+1} =& g + \sigma \varepsilon^c_{t+1} + (J_{t+1} - \lambda\omega) \\ \Delta d_{t+1} =& g + \varphi\sigma\varepsilon^d_{t+1} + \varphi\,(J_{t+1} - \lambda\omega)\,, \end{aligned} \tag{12}$$

where ε^c and ε^d are bivariate standard normal with correlation ϱ, and J equals a Poisson mixtures of normals,

$$\begin{aligned} x \sim& \text{Poisson}(\lambda) \\ J_{t+1}|x \sim& N(x\omega, x\nu^2). \end{aligned} \tag{13}$$

Both the number of jumps (x) and the jump size are IID over time and independent of $(\varepsilon^c, \varepsilon^d)$. Because (12) demeans jumps with their expected value, $\lambda\omega$, g represents the mean of Δc and Δd. The values of g, φ, ϱ, and β are kept unchanged relative to Section 4.1. To keep the unconditional consumption volatility fixed at 3% per year, we set $\sigma = [0.03 - \lambda(\omega^2 + \nu^2)]/\sqrt{12}$ (see Backus et al. 2011, equation 6). We characterize the quantity of tail risk in monthly returns by their skewness (-0.685), kurtosis (7.22), and probability of falling between -34% and -11.3% (0.011), and replicate these moments in the model by setting $\omega = -2.09\%$, $\nu = 2.108\%$, and $\lambda = 0.26164$ (one jump every 3.8 years on average). Jumps are therefore considerably smaller and more likely than in typical calibrations of the rare disaster model. As in Section 4.1, we choose the coefficient of relative risk aversion, $\gamma = 14.89$, to generate an equity premium of 8% per year. We refer to this calibration as "Model 2".

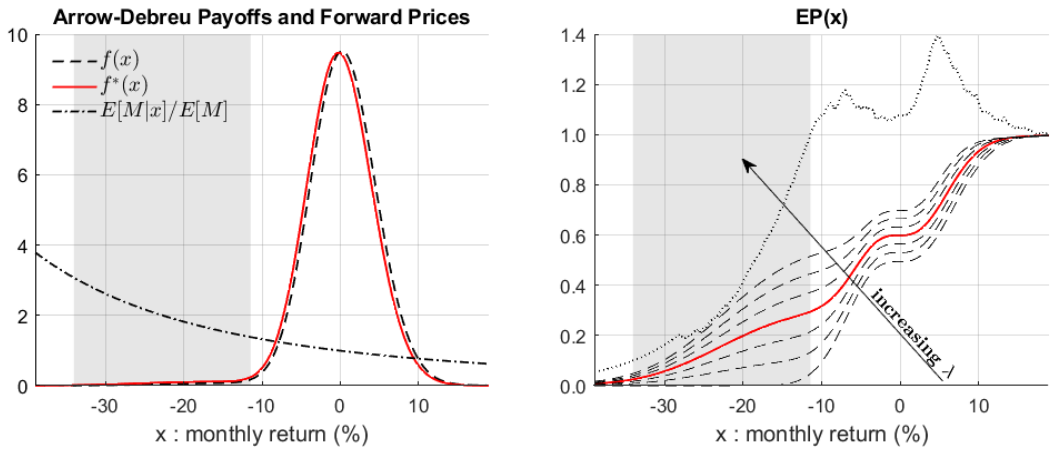

Figure V: Left-skewed shocks (Model 2).

Figure V shows average prices and payoffs of Arrow-Debreu securities in the left panel, and the $EP(x)$ curve they imply in the right panel (solid line). The projected pricing kernel is almost unchanged relative to the lognormal model and, as a result, f^* once again resembles a left-shifted version of f. Different from the lognormal case, however, both $EP(x)$ and the two densities remain positive until far into the left tail. Table II shows that returns in the gray-shaded region account for 27/100 of the equity premium – almost six times their contribution in the lognormal model, but still far below the 9/10 contribution in the data.

To investigate how sensitive this finding is with respect to the exact quantity of tail risk, we consider calibrations with different jump intensities (λ). These calibrations are illustrated with dashed lines in the right panel of Figure V. For each value of λ, risk aversion is chosen to generate an equity premium of 8% per year, while other parameters are left unchanged. The lowest line corresponds to the lognormal calibration of Model 1 ($\lambda = 0$), whereas the highest line doubles the frequency of jumps relative to the calibration of Model 2 ($\lambda = 0.26164 \times 2$). In the later case, returns have a skewness of -1.72, kurtosis of 12.38, and probability of 0.019 of falling in the gray-shaded region – an amount of tail risk that far exceeds what we observe

in the data. Nevertheless, we find that states in the gray-shaded region account for only 47/100 of the equity premium. For comparison, we also compute our metrics for the calibration of Backus et al. (2011, column 4 of their Table II), which assumes somewhat smaller but more frequent jumps. We find that their calibration implies that returns between -34% and -11.3% have a probability of 0.021 and account for 30.4/100 of the equity premium. This evidence suggests that a larger (and realistic) quantity of tail risk than implied by lognormality is insufficient for capturing the importance of left tail events.

5.1 Habits and long-run risks with skewed shocks

We consider the models of Bekaert and Engstrom (2017) and Drechsler and Yaron (2011) to evaluate whether the implications of Model 2 carry over to environments with richer dynamics. Bekaert and Engstrom (2017) augment the Campbell-Cochrane model with left-skewed shocks in consumption and dividend growth rates. Table II shows that their calibration implies that states between -34% and -11.3% have a probability of 0.028 (versus 0.011 in the data) and contribute 36.9/100 to the equity premium. Drechsler and Yaron (2011) add jumps to the state variable processes of the Bansal-Yaron model. We find that their calibration implies a probability of 0.017 and equity premium contribution of 35.7/100 for the same states. For both models, the price of tail risk is similar to the one in the original calibrations. Despite implying a quantity of tail risk that exceeds its empirical counterpart, however, the Bekaert and Engstrom (2017) and Drechsler and Yaron (2011) models attribute a significantly smaller role to states between -34% and -11.3% than the data. The states' empirical equity premium contribution of 9/10 has a small sample p-value below 5% in both models.

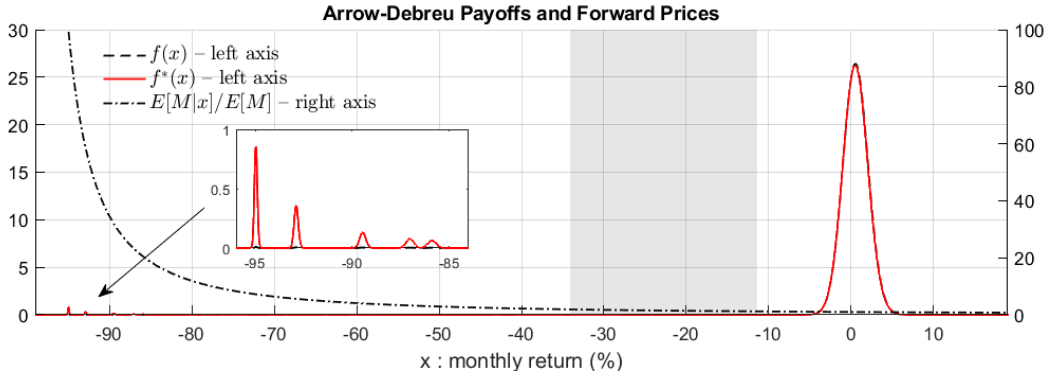

Figure VI: The Barro (2009) model.

5.2 Rare disasters

Another perspective on the quantity of tail risk comes from the rare disaster literature initiated by Rietz (1988). Barro (2006, 2009) measures large consumption declines in international data and shows that they imply a realistic equity premium in an IID representative agent model with a risk aversion coefficient of 4. His calibration can be interpreted as an extreme version of Model 2, where jumps occur rarely but take on disastrous magnitudes. Barro's main results are based on an unlevered consumption claim. We assume a leverage parameter of 2.6, the value used by Wachter (2013). Figure I shows that the model attributes more than half of the equity premium to returns below -89%, but very little to states in the gray-shaded region. Figure VI shows average forward prices and payoffs of Arrow-Debreu securities for the model. The low degree of risk aversion results in a flat pricing kernel for small to moderately large negative returns, so that f and f^* look nearly identical in that region. In contrast, the pricing kernel rises to a value of about 100 in the far left tail, and Arrow-Debreu forward prices exceed the securities' expected payoffs by two orders of magnitude. The single worst disaster realization alone (out of 87 possible disasters in the calibration) accounts for about 1/3 of the equity premium.

Wachter (2013) augments the Rietz-Barro model with a long-run risks channel by assuming persistent variation in the probability of disasters and an Epstein-Zin agent with a preference for the early resolution of uncertainty. Figure I shows that her calibration attributes about 2/3 of the equity premium to "disaster states" in the far left tail and 1/3 to "long-run risk states" to the right of the gray-shaded region. Despite a probability of 0.007, however, states in the gray-shaded region account for only 4.7/100 of the equity premium. The states' empirical equity premium contribution of 9/10 has a small sample p-value below 1%. Its reliance on extremely negative shocks therefore makes the rare disaster mechanism inconsistent with the importance of more moderate negative return states, regardless of whether or not the probability of disasters varies over time.

6 Increasing the price of tail risk

The price of risk for different return states reflects the relationship between returns and consumption and the relationship between consumption and marginal utility. To increase the price of risk for tail events in the stock market, a model therefore either has to make large, adverse shocks to dividends relatively more likely to coincide with large, adverse shocks to consumption, or increase the level of marginal utility associated with adverse shocks to consumption. An important challenge in this context is that we cannot simultaneously increase the price of risk for small negative returns, because doing so would result in an unrealistically large equity premium. Simply increasing risk aversion or making all shocks more correlated can therefore not increase the price of tail risk. We discuss two mechanisms that overcome this tension and illustrate their quantitative implications in simple IID settings.

6.1 Co-tail risk

The models considered so far exclusively rely on two correlation structures for consumption and dividend growth. The first approach assumes that consumption and dividends are perfectly correlated, $\Delta d = \varphi \Delta c$, as in Abel (1999). The IID models in Sections 4.1 and 5 and the models of Barro (2009), Backus et al. (2011), Wachter (2013), and Bekaert and Engstrom (2017) fall in this category. The second approach, exemplified by Campbell and Cochrane (1999), Bansal and Yaron (2004), and Drechsler and Yaron (2011), models shocks to consumption and dividends as (conditionally) jointly normal. Because the bivariate normal distribution relies on a single correlation coefficient, an increase in the dependence of large shocks necessitates a simultaneous increase in the dependence of small shocks. Thus, neither approach makes it possible to increase the *relative* importance of large dividend shocks for the equity premium.

An alternative correlation structure results if large negative shocks to consumption and dividends are likely to coincide, while other shocks are not. We illustrate this mechanism based on the Poisson-normal endowment in Equations (12) and (13) by assuming that jumps are common (as before) and Gaussian shocks are imperfectly correlated. Starting with the calibration in Section 5 ($\varrho = 1$), which corresponds to the lowest $EP(x)$ curve in Figure VII, we gradually lower the correlation of Gaussian shocks and raise the coefficient of relative risk aversion to keep the equity premium fixed at 8%. All other parameters are kept unchanged relative to Section 5. Because the marginal shock distributions are not altered in these calibrations, the process continues to match the quantity of tail risk in monthly returns.

As the correlation of Gaussian shocks is lowered and risk aversion is raised, jumps become increasingly more important for the equity premium. The highest $EP(x)$ curve in Figure VII corresponds to a calibration with independent Gaussian

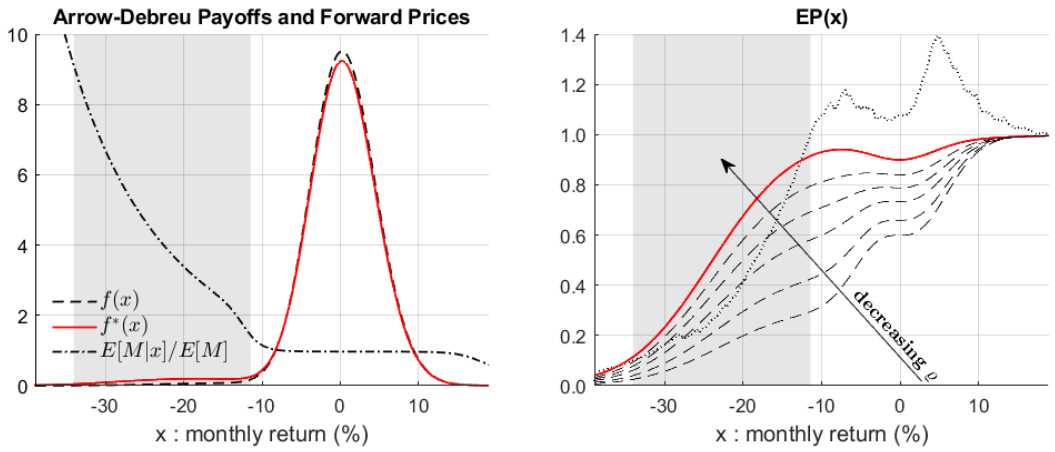

Figure VII: Increased co-tail risk (Model 3).

shocks ($\varrho = 0$; henceforth "Model 3"), for which the left panel shows average state prices and payoffs. Relative to the perfect correlation case, the model's implications change in several important ways. First, the pricing kernel is approximately horizontal and below average ($\mathbb{E}[M|R] < \mathbb{E}[M]$) for small absolute returns. As a result, $EP(x)$ is decreasing for small negative returns and increasing for small positive returns, as in the data (recall the discussion in Section 2.2). Second, the pricing kernel is considerably steeper for states in the gray-shaded region, resulting in a steep $EP(x)$ curve and an equity premium contribution of 80.1/100 for the associated states. The (forward price)-to-(average payoff) ratio of an Arrow-Debreu security that pays \$1 in the gray-shaded region, our proxy for the price of tail risk, increases to 3.13, very close to the empirical value of 3.24 (see Table II). In fact, the entire $\mathbb{E}[M|x]/\mathbb{E}[M]$ curve strongly resembles its empirical counterpart in Figure II (recall that $\mathbb{E}[M|x]/\mathbb{E}[M] = f^*/f$ in IID models). Third, f^* no longer looks like a left-shifted version of f. Arrow-Debreu forward prices significantly exceed the securities' expected payoffs in the far left tail, but fall below them for states to the right of the gray-shaded region, similar to the empirical densities in Figure II. Table II further shows that $EP(x)$, the quantity of risk, and the price of risk for states in

the gray-shaded region are not significantly different from the data in finite samples. Consistent with the main channel of Model 3, Schreindorfer (2019) shows empirically that the "downside correlation" between consumption and dividend growth is increasing in the left tail, i.e. large negative shocks are more likely to coincide.

6.2 Tail-sensitive risk preferences

Standard risk preferences restrict the relationship between the marginal utility for shocks of different magnitudes in a way that resembles how the bivariate normal distribution restricts the dependence between different shocks. In particular, all models considered above assume utility functions that rely on a single risk aversion coefficient. As such, they can only increase the marginal utility of large shocks by simultaneously increasing the marginal utility of small shocks.

One way to overcome this tension is to assume a utility function that disentangles risk attitudes towards small and large risks, such as Generalized Disappointment Aversion (GDA) (Gul 1991; Routledge and Zin 2010). GDA is an axiomatic extension of expected utility theory that nests expected utility as a special case and rationalizes the Allais (1979) paradox. When GDA is used to govern risk preferences in the intertemporal utility function of Epstein and Zin (1989), the pricing kernel in IID environments equals

$$M_{t+1} = \tilde{\beta} \times \underbrace{e^{-\gamma \Delta c_{t+1}}}_{\text{small risks}} \times \underbrace{\left(1 + \theta \mathbf{1}\{\Delta c_{t+1} \leq \ln(\delta) + \tilde{x}\}\right)}_{\text{tail risks}}. \tag{14}$$

$\tilde{\beta}$ and $\tilde{x}$ are endogenous constants (see Schreindorfer (2019) for details). For $\theta = 0$ (no disappointment aversion) and an elasticity of intertemporal substitution (EIS) of $1/\gamma$ (time separable utility), the pricing kernel equals the one of Models 1, 2, and 3. For $\theta > 0$, the "tail risks term" overweights consumption growth realizations in the left tail. The parameter θ governs the magnitude of the disappointment penalty

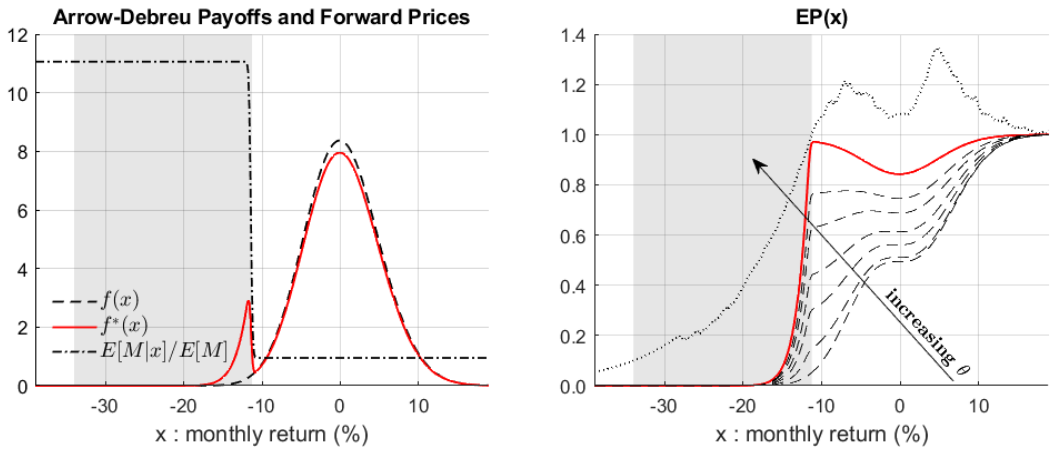

Figure VIII: Tail-sensitive risk preferences (Model 4).

and δ governs the threshold below which outcomes are considered disappointing.

To illustrate the effect of GDA on risk premia in the clearest possible way, we consider an IID lognormal endowment and perfectly correlated shocks. The lowest $EP(x)$ curve in Figure VIII corresponds to Model 1 (time separable CRRA utility). Starting with this calibration (for which $\theta = 0$) and a disappointment threshold of $\delta = 0.987$, we lower γ and increase θ to keep the equity premium fixed at 8%. All other parameters are unchanged relative to Model 1. As γ is lowered and θ is raised, adverse consumption shocks become more disappointing and states in the gray-shaded region become more important for the equity premium. The highest $EP(x)$ curve corresponds to a calibration with $\gamma = 0$ and $\theta = 10.64$ (henceforth "Model 4"), for which the left panel shows average state prices and payoffs. States in the gray-shaded region coincide with a pricing kernel above 10, a steep $EP(x)$ curve, a price of risk of 9.935, and they account for 96.7/100 of the equity premium (see Table II). Clearly, GDA risk preferences allow for substantially more aversion against extreme events than standard utility functions and can therefore help rationalize the importance of such states for the equity premium.

Two features make Model 4 very stylized. First, the left tail of the lognormal

distribution approaches zero at a very fast rate, which implies that probability mass in the gray-shaded region is concentrated around its right endpoint. Such states are therefore the primary source of risk premia and the $EP(x)$ curve rises abruptly, rather than gradually, as in the data. This feature can be eliminated by assuming endowment shocks from a left-skewed distribution that increases the relative likelihood of returns in the far left tail. Second, the projected pricing kernel and Arrow-Debreu forward prices have a discontinuity at the disappointment threshold. This feature can be eliminated by specifying imperfectly correlated shocks to consumption and dividends, so that returns (which reflect dividend shocks) become gradually more likely to coincide with disappointments (which reflect consumption shocks) when they fall further into the left tail.

6.3 Schreindorfer (2019)

Schreindorfer (2019) combines the channels of Models 3 and 4. His GDA model assumes that consumption and dividends are subject to independent Gaussian shocks and negatively exposed to a common exponential shock. In contrast to Model 3, where large negative shocks (jumps) always coincide and small shocks (Gaussian) never coincide, his setting implies that shocks become gradually more likely to coincide the more negative they are. Schreindorfer (2019) calibrates the model to replicate the average returns and Sharpe ratios of options with different strikes, but we find that it provides a similarly good match for our equity premium decomposition. States in the gray-shaded region contribute 72/100 to the equity premium (versus 90/100 in the data), occur with a probability of 0.008 (versus 0.011 in the data), and coincide with a price of risk of 4.69 (versus 3.24 in the data).

7 Last thoughts

The asset pricing literature has proposed a number of modifications to utility functions and shock structures to overcome the empirical challenges of Lucas's (1978) consumption-based asset pricing framework. The models of Campbell and Cochrane (1999), Bansal and Yaron (2004), and Rietz (1988) and Barro (2006, 2009) all account for the average equity premium and related asset market facts based on plausible macroeconomic risks. Due to their quantitative success and economic appeal, the same mechanisms have been adopted to understand a wide range of issues linked to risk premia, including exchange rate movements (Verdelhan 2010; Colacito and Croce 2011; Farhi and Gabaix 2015), the term structure of interest rates (Wachter 2006; Bansal and Shaliastovich 2013), credit spreads (Chen 2010; Bhamra et al. 2010), business cycles (Gourio 2012; Croce 2014), labor markets (Kilic and Wachter 2018; Favilukis et al. 2019), and climate change policies (Bansal et al. 2016).

We document empirically that stock market risk premia predominantly compensate investors for shocks that coincide with infrequent, large, but not extremely large negative returns. This fact is quantitatively at odds with leading calibrations of the aforementioned theories, including those that generate a realistic quantity of tail risk in returns.

Our evidence points to a higher price of tail risk than implied by existing models, and we propose two ways to generate one. One option is to assume that large negative shocks to equity markets are likely to coincide with large drops in consumption, while smaller shocks to stock prices are only weakly correlated with the real economy. We model this channel via shocks to consumption and dividends, but it could similarly be applied to shocks to economically-important state variables, such as the surplus consumption ratio in the habit model or conditional consumption growth moments in the long-run risks model. Another option is to postulate a utility func-

tion that disentangles risk attitudes towards small and large risk. We illustrate this channel via GDA risk preferences, which are straightforward to embed into any model with Epstein and Zin (1989) utility. We show that both mechanisms can replicate key features of our equity premium decomposition in IID representative agent models. An important route for future research is to investigate whether the same mechanisms can align existing asset pricing models more closely with sources of the equity premium, while also preserving their implications for excess volatility and the predictability of stock market returns. More broadly, we believe that future generations of asset pricing models, representative agent or not, should pay close attention to the composition of risk premia, in addition to their level. $EP(x)$ provides a tool for doing so.

A Appendix

A.1 Data and estimation

We "clean" the option dataset by remoing observations if they (i) violate the static no-arbitrage bounds $P \leq K/R^f$ or $C \leq S$, (ii) have a best bid quote of zero, (iii) have the CBOE's error code 999 for ask quotes or 998 for bid quotes, (iv) have non-positive bid-ask spreads, (v) have midquotes less than \$0.50, (vi) are singles (a call quote without a matching put quote or vice versa), (vii) are PM settled, or (viii) have annualized implied volatilities (IVs) less than 2% or more than 200%. To detect any additional "outliers", we fit a simple linear function in maturity and standardized moneyness (see definitions below) to IVs on each date, and remove observations that are highly influential based on their Cook's distance (a common statistical metric for detecting outliers).

We use the Black and Scholes (1973) formula to convert option prices to implied volatility (IV) units, fit an interpolant to them, and then map interpolated IVs back to prices. The benefit of doing so is that IVs can easily be fit with a parsimonious function due to their homogeneous magnitudes, whereas option prices can not. Importantly, however, this approach does not assume the validity of Black and Scholes's assumptions because their formula is merely used to map back and forth between two spaces.

For a given maturity, we parameterize IVs with the function

$$IV(\kappa) = \left[a + b\left(\rho(\kappa - m) + \phi(\kappa - m)^2 + \sqrt{(\kappa - m)^2 + \sigma^2}\right)\right]^{\theta}, \tag{A.1}$$

where $\kappa = \frac{\ln(K/S_t)}{VIX_t\sqrt{T}}$ denotes the option's moneyness in standard deviation units (henceforth "standardized moneyness") and T denotes its maturity in years. For example, a put option with $\kappa = -2$ and $T = 1/2$ is two standard deviations out-of-the-money and has 6 months to maturity. The θ-exponent determines whether

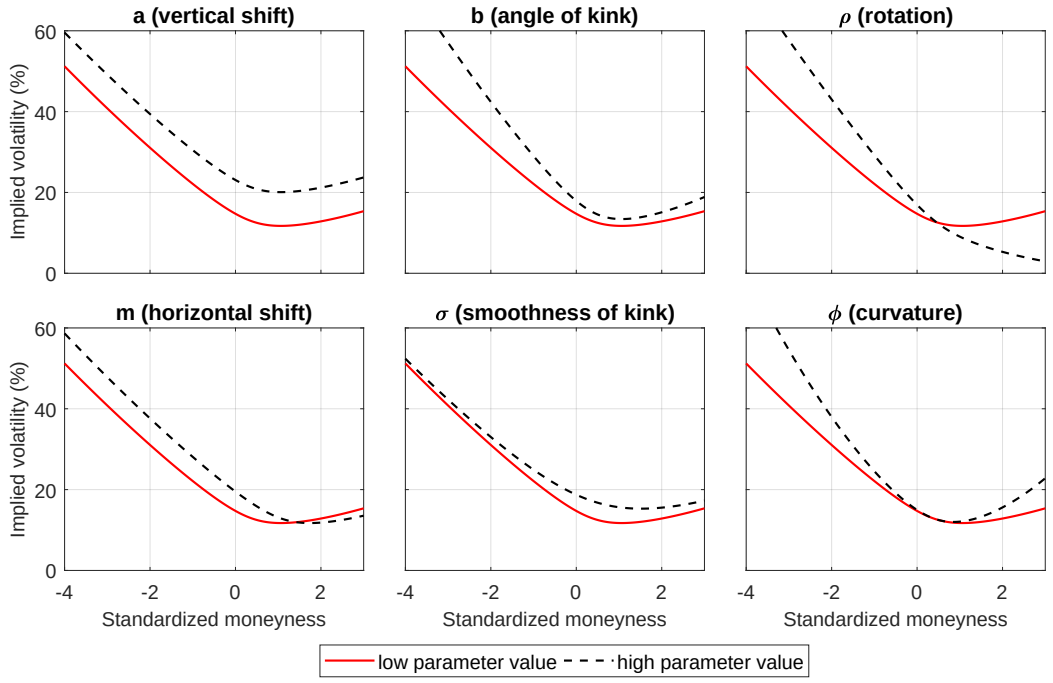

Figure IX: Comparative statics for the implied volatility parameterization. The functional form is shown in Equation A.1; θ is set to 1.

the function applies to volatility ($\theta = 1$) or variance ($\theta = 1/2$) units, which in turn affects the tail behavior of the IV curve. The remaining six parameters give the $IV(\kappa)$ function the flexibility to take on a variety of shapes, as illustrated in Figure IX. Beyond smoothness, $IV(\kappa)$ therefore imposes very little structure on the shape of the IV curve. Equation (A.1) represents a slight generalization of the "SVI parameterization", a method widely used by practitioners to fit IV curves, see, e.g. Gatheral (2006).

Because we would like to interpolate option prices in both the moneyness and maturity dimension, we extend (A.1) in the maturity dimension by specifying linear functions of T for the coefficients, e.g.,

$$a = a_0 + a_1 T, \tag{A.2}$$

and similarly for $(b, \rho, m, \phi, \sigma)$. The linear functional form is chosen for its parsimony – it results in twelve free parameters for the bivariate function $IV(\kappa, T)$ – and to impose the economically-motivated constraint that $IV(\kappa)$ changes smoothly as a function of maturity.

For each day in the sample, we find the parameter vector $\theta \equiv (a_j, b_j, \rho_j, m_j, \phi_j, \sigma_j)$, $j \in \{0, 1\}$, by minimizing a two-part objective function, $O_t(\theta) = O_{1,t}(\theta) + O_{2,t}(\theta)$. The first term is given by the root mean squared error between observed IVs and the function $IV(\kappa, T; \theta)$,

$$O_{1,t}(\theta) = \sqrt{\frac{1}{N_t} \sum_{i=1}^{N_t} [IV_i - IV(\kappa_i, T_i; \theta)]^2}, \tag{A.3}$$

where N_t is the number of observations on day t. The second term penalizes parameter vectors for which the $IV(\kappa, T; \theta)$-implied density of forward state prices does not integrate to one. To compute it, we evaluate $IV(\kappa, T; \theta)$ on a fine grid of standardized moneyness between -10 and +3 standard deviations and maturities between 10 to 120 days, map IVs back to put prices based on the Black-Scholes formula, and compute forward state prices for each maturity via finite differences based on Equation (10). The penalty is set to

$$O_{2,t}(\theta) = 2 \sum_{T} \times \left| \int f_t^*(\kappa, T; \theta) d\kappa - 1 \right|, \tag{A.4}$$

where $f_t^*(\kappa, T; \theta)$ are forward state prices for maturity T on day t, and we sum over all values in the maturity grid.

We use a particle swarm algorithm to minimize $O_t(\theta)$ globally over the parameter space and discard any set of parameters for which option prices on the grid are not positive, monotone, and convex. We include observations of put options with a standardized moneyness below 0.5, call options with a standardized moneyness above -0.5, and maturities between 8 and 120 days. In-the-money options and long-

Table II: Alternative state price estimates

	R^2		f^* moments			Return (in %) at which $EP(x) =$			
	mean	median	vol	skew	kurt	0.05	0.10	0.50	1.00
(A)	0.930	0.973	0.191	-1.56	10.72	-39.7	-34.1	-18.4	-11.5
(B)	0.919	0.961	0.191	-1.54	10.51	-39.6	-34.0	-17.9	-11.3
(C)	0.932	0.977	0.191	-1.57	10.85	-40.1	-34.3	-18.3	-11.2
(D)	0.886	0.930	0.189	-1.39	10.10	-39.4	-33.1	-16.1	-8.5

Statistics for alternative estimates of 30-day forward state prices. The estimates differ by the restrictions they impose on Equation A.1: (A) $\theta = \frac{1}{2}$, (B) $\theta = \frac{1}{2}, \phi = 0$, (C) $\theta = 1$, and (D) $\theta = 1, \phi = 0$.

maturity options are excluded because they have substantially larger bid-ask spreads and trade less frequently. Contracts in the week before expiration are also excluded.

Due to its flexibility, it is not surprising that the parameterized IV surface provides a good fit to observed IVs. Outside of the observed moneyness region, however, $IV(\kappa, T)$ is only restricted by theoretical no arbitrage constraints, and could therefore be sensitive to the parametric functional form we have chosen. To evaluate this possibility, we estimate four versions of Equation (A.1) that differ in terms of their behavior for extreme moneyness values. We consider either variance ($\theta = 1/2$) or volatility ($\theta = 1$) units, and either restrict the quadratic term in Equations (A.1) to zero ($\phi = 0$) or not ($\phi \neq 0$). For all four cases, Table II shows implied R^2's for the fit to observed IVs, higher moments of the implied 30-day f^* density, and return thresholds at which the resulting 30-day $EP(x)$ curve equals 0.05, 0.1, 0.5, or 1.

Clearly, specification D (volatility units, no quadratic term), for which both "legs" of the IV curve are linear by construction, provides the worst fit. The average R^2 across the 7,304 trading days in our sample equals 0.886, whereas the median R^2 equals 0.93. Specifications A-C, all of which allow for concavity in the legs of the

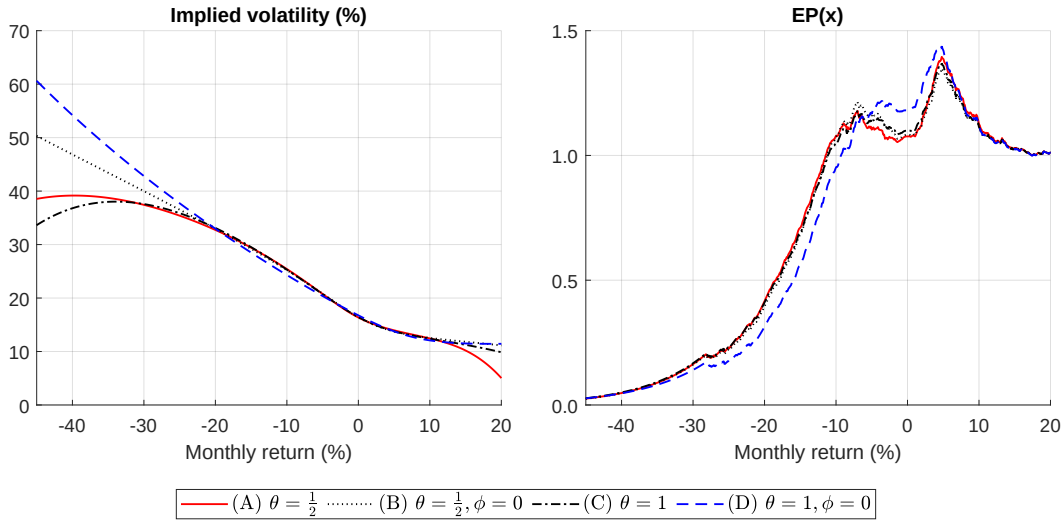

Figure X: Alternative estimates of implied volatilities and $EP(x)$.

IV curve, provide better and very similar fits: The average R^2 falls between 0.919 and 0.932, whereas the median R^2 falls between 0.961 and 0.977. The left panel of Figure X further shows that the average 30-day IV curve looks nearly identical for returns between -20% and +10% (which contains most option observations) in these three cases, but noticeably different in both tails (which contains few observations). Despite these differences, however, Table II and the right panel of Figure X show that the higher moments of f^* and the shape of the resulting $EP(x)$ curve are nearly identical across specifications A-C. Our main results are therefore robust to the way IVs are extrapolated to the tails. The main part of the paper relies on the $f^*(R)$ estimate from specification B (variance units, no quadratic term) because it represents the version of Equation (A.1) suggested by Gatheral (2006).

A.2 Computation

Apart from the Drechsler and Yaron (2011) model, which is high-dimensional and solved via log-linearization, we solve all models using globally accurate numerical

solution methods. The online appendix provides details. Based on a given model solution, we compute $f(R)$ via a Monte Carlo simulation. To begin, we construct an equally-spaced grid $r_1, \ldots, r_N$ for returns. Denote the bin width by $\delta = r_{i+1} - r_i$ and the i^{th} bin midpoint by $\bar{r_i} = \frac{r_i + r_{i+1}}{2}$. Next, simulate R_{t+1} for T periods.[6] The value of the physical distribution at the i-th grid midpoint is computed as

$$f\left(\bar{r_i}\right) = \frac{1}{\delta \times T} \sum_{t=1}^{T} \mathbf{1} \left\{ r_i < \frac{S_{t+1}}{S_t} \leq r_{i+1} \right\}, \quad i = 1, ..., N-1 \tag{A.5}$$

where $\mathbf{1}\{\cdot\}$ is an indicator and $N = 10,000$. We set the simulation length to $T = 10^{11}$ periods, which reduces Monte Carlo noise to essentially zero for all models.

To find $f^*(R)$, we compute state-contingent put option prices on the aforementioned return grid and a fine grid for each model's state. Next, we find the state-contingent risk-neutral density via finite differences based on Equation 10, simulate the state for a large number of periods, and compute the average density as

$$f^*(\bar{r_i}) = \frac{1}{L} \sum_{l=1}^{L} f^*(\bar{r_i}; X_l), \quad i = 1, ..., N-1 \tag{A.6}$$

where X_l denotes the simulated state in period l. $EP(x)$ is given by

$$EP(x) = \frac{\sum_{i=1}^{N} \mathbf{1}\left\{r_i \leq x\right\} r_i (f(r_i) - f^*(r_i))\delta}{\sum_{j=1}^{N} r_j (f(r_j) - f^*(r_j))\delta}. \tag{A.7}$$

For each model we consider, the online appendix discusses implementation details for $EP(x)$, as well as results from the original studies that we replicated to verify the accuracy of our numerical solutions.

[6] Steps for doing so are model specific, but generally involve simulating shocks, using them to evaluate dividend growth and paths of the model's state variable(s), using the state variable(s) to evaluate the price-dividend ratio, and computing $R_{t+1} = \frac{PD_{t+1}}{PD_t} e^{\Delta d_{t+1}}$.

References

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

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.

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

Bansal, R., M. Ochoa, and D. Kiku (2016): "Climate Change and Growth Risks," Working Paper.

Bansal, R. and I. Shaliastovich (2013): "A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets," *Review of Financial Studies*, 26, 1–33.

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.

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

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, 645–703.

Black, F. and M. Scholes (1973): "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637–654.

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

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.

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.

Carr, P. and L. Wu (2009): "Variance Risk Premiums," *Review of Financial Studies*, 22, 1311–1341.

Chen, H. (2010): "Macroeconomic Conditions and the Puzzles of Credit Spreads and Capital Structure," *Journal of Finance*, 65, 2171–2212.

Colacito, R. and M. M. Croce (2011): "Risks for the Long Run and the Real Exchange Rate," *Journal of Political Economy*, 119, 153–181.

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

Croce, M. M. (2014): "Long-run productivity risk: A new hope for production-based asset pricing?" *Journal of Monetary Economics*, 66, 13–31.

Cuesdeanu, H. and J. Jackwerth (2018): "The pricing kernel puzzle: survey and outlook," *Annals of Finance*, 14, 289–329.

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. and A. Yaron (2011): "What's Vol Got to Do with It," *Review of Financial Studies*, 24, 1–45.

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.

Farhi, E. and T. Gabaix, Xavier (2015): "Rare Disasters and Exchange Rates," *Quarterly Journal of Economics*, 131, 1–52.

Favilukis, J., X. Lin, and X. Zhao (2019): "The Elephant in the Room: the Impact of Labor Obligations on Credit Risk," *American Economic Review*, forthcoming.

Figlewski, S. (2008): "Estimating the Implied Risk Neutral Density for the U.S. Market Portfolio," in *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*, ed. by T. Bollerslev, J. R. Russell, and M. Watson, Oxford University Press.

Gatheral, J. (2006): *The Volatility Surface: A Practitioner's Guide*, Wiley Finance.

Gourio, F. (2012): "Disaster risk and business cycles," *American Economic Review*, 102, 2734–66.

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

Kilic, M. and J. A. Wachter (2018): "Risk, unemployment, and the stock market: A rare-event-based explanation of labor market volatility," 31, 4762–4814.

Lucas, R. E. (1978): "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1429–1445.

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

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

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

Rubinstein, M. (1994): "Implied Binomial Trees," *Journal of Finance*, 49, 771–818.

Schreindorfer, D. (2019): "Macroeconomic Tail Risks and Asset Prices," *Review of Financial Studies*, forthcoming.

Schulz, F. (2016): "On the Timing and Pricing of Dividends: Comment," *American Economic Review*, 106, 3185–3223.

Seo, S. B. and J. A. Wachter (2018): "Do Rare Events Explain CDX Tranche Spreads?" *Journal of Finance*, 73, 2343–2383.

van Binsbergen, J., M. Brandt, and R. Koijen (2012): "On the Timing and Pricing of Dividends," *American Economic Review*, 102, 1596–1618.

Verdelhan, A. (2010): "A Habit-Based Explanation of the Exchange Rate Risk Premium," *Journal of Finance*, 65, 123–146.

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

Wachter, J. A. (2006): "A consumption-based model of the term structure of interest rates," *Journal of Financial Economics*, 79, 365–399.

Welch, I. (2018): "The (Time-Varying) Importance of Disaster Risk," *Financial Analysts Journal*, 72, 14–30.