

A Tale of Two Tails: A Model-free Approach to Estimating Disaster Risk Premia and Testing Asset Pricing Models

Tjeerd de Vries*

January 6, 2024

[CLICK HERE FOR THE MOST RECENT VERSION](#)

Abstract

I introduce a model-free methodology to assess the impact of disaster risk on the market return. Using S&P500 returns and the risk-neutral quantile function derived from option prices, I employ quantile regression to estimate local differences between the conditional physical and risk-neutral distributions. The results indicate substantial disparities primarily in the left-tail, reflecting the influence of disaster risk on the equity premium. These differences vary over time and persist beyond crisis periods. On average, the bottom 5% of ex-ante returns contribute to 17% of the equity premium, shedding light on the Peso problem. I also find that disaster risk increases the stochastic discount factor's volatility. Using a lower bound observed from option prices on the left-tail difference between the physical and risk-neutral quantile functions, I obtain similar results, reinforcing the robustness of my findings.

Keywords: Asset pricing, Disaster Risk, Quantile methods

JEL Codes: G13, G17, C14, C22

*Department of Economics, University of California San Diego. Email: tjdevrie@ucsd.edu. I am particularly grateful to my advisors, Allan Timmermann and Alexis Toda, for their support and feedback throughout the writing process. I would also like to thank James Hamilton, Xinwei Ma, Yixiao Sun, Rossen Valkanov, Brendan Beare, Caio Almeida, Ali Uppal, Rob Engle (discussant), Andrew Lo (discussant), Lerby Ergun, Casper de Vries, Jens Jackwerth, Stefan Faridani and seminar participants of the 2022 SoFiE conference, 2023 JOIM conference and ESI Brown Bag Seminar for useful feedback. All errors are my own.

1 Introduction

Disaster risk has emerged as a pervasive and influential concept in asset pricing, offering a prominent explanation of the equity premium puzzle, as well as other asset pricing puzzles.¹ Little is known, however, about the quantitative properties of disaster risk and evidence for it is often inferred indirectly, such as from the historically high equity premium. Nevertheless, a high equity premium does not necessarily arise due to disaster risk, and the literature has yet to reach an unambiguous conclusion regarding its ability to explain asset pricing puzzles (see, e.g., [Julliard and Ghosh \(2012\)](#)). [Ross \(2015\)](#) refers to disaster risk as dark matter and summarizes the concept as follows: “It is unseen and not directly observable but it exerts a force that can change over time and that can profoundly influence markets”.

In this paper, I propose a model-free methodology to measure and track disaster risk in S&P500 returns through time. My results unequivocally show that disaster risk is pervasive and is a primary determinant of the equity premium. In establishing these results, I confront two critical challenges that have hindered inference so far about disaster risk. Firstly, to estimate disaster risk in a model-free manner, the literature often estimates the stochastic discount factor (SDF), defined as the ratio of risk-neutral to physical density. Disaster risk is then thought of as the tendency of the SDF to take large values in the left-tail of the return distribution. However, this approach faces scrutiny due to the potential for erratic results when estimating the density ratio in the tails, thereby complicating robust inference. Secondly, it is crucial to account for changing conditioning information. Typically, estimation of the physical density involves pooling historical returns, while the risk-neutral density relies on forward-looking option prices. This disparity in information sets can lead to inconsistent estimates of conditional disaster risk.

To address these challenges, I consider an approach that avoids the need for density estimation. Starting from the absence of arbitrage opportunities, a risk-neutral distribution exists that can be identified from option prices without assuming any model ([Breedon and Litzenberger, 1978](#)). However, the conditional physical distribution, which describes the actual evolution of the market return, remains unobserved. To proceed, I use quantile regression (QR) to estimate

$$\underbrace{Q_{t,\tau}(R_{m,t \rightarrow N})}_{\text{Unobserved}} = \beta_0(\tau) + \beta_1(\tau) \underbrace{\tilde{Q}_{t,\tau}(R_{m,t \rightarrow N})}_{\text{Observed}} \quad \tau \in (0, 1), \quad (1.1)$$

¹See, for example, [Rietz \(1988\)](#), [Barro \(2006, 2009\)](#), [Drechsler and Yaron \(2011\)](#), [Gabaix \(2012\)](#), [Wachter \(2013\)](#), [Constantinides and Ghosh \(2017\)](#), [Isoré and Szczerbowicz \(2017\)](#), [Farhi and Gourio \(2018\)](#), [Seo and Wachter \(2019\)](#) and [Schreindorfer \(2020\)](#).

where $Q_{t,\tau}$ and $\tilde{Q}_{t,\tau}$ represent the physical and risk-neutral τ -quantiles, respectively, of the market return $R_{m,t \rightarrow N}$, from period t to $t + N$. The parameters in (1.1) can be estimated using quantile regression, with the observed time series of returns, $\{R_{m,t \rightarrow N}\}_{t=1}^T$, as the dependent variable and $\{\tilde{Q}_{t,\tau}\}_{t=1}^T$ as the regressor. Importantly, both $R_{m,t \rightarrow N}$ and $\tilde{Q}_{t,\tau}$ are conditioned on the same information set.

In general, quantile regression estimates the best linear approximation to the physical quantile function, from which $R_{m,t \rightarrow N}$ is drawn. But because the risk-neutral quantile function is a highly non-linear transformation of state variables, the estimation accommodates non-linear dependence between the physical quantile function and the state variables it depends on. Hence, any deviation from the risk-neutral benchmark, $[\beta_0(\tau), \beta_1(\tau)] = [0, 1]$, signifies a local difference between the physical and risk-neutral measures at the τ -quantile. Since the equity premium is determined by these differences, it is natural to define *disaster risk premia* as the difference between $Q_{t,\tau}$ and $\tilde{Q}_{t,\tau}$ in the left-tail, i.e. for values of τ close to zero.

Based on the QR estimates, two key findings emerge: (i) the risk-neutral benchmark cannot be rejected in the right-tail ($\tau \geq 0.7$) but it is rejected in the left-tail ($\tau \leq 0.3$); and (ii) the in-sample and out-of-sample explanatory power of the risk-neutral quantile is significantly higher in the right-tail compared to the left-tail. Both findings suggest that disaster risk is the main driver of the equity premium.

Building on these results, I estimate the conditional Lorenz curve and Gini coefficient associated with the equity premium. These statistics summarize how much the conditional equity premium is driven by the lowest returns (disaster), akin to its interpretation of wealth inequality in labor economics. I find that the Lorenz curve is always concave, and the Gini coefficients are far above zero in every time period, thus showing that disaster risk is a pervasive feature of the data. On average, I find that ex-ante returns below the 5th percentile contribute to 17% of the total equity premium.

While this result demonstrates that disaster risk is an important driver of expected returns, it also adds nuance to the degree of disaster risk necessary to explain the equity premium. In particular, previous papers attribute about 90% of the equity premium to returns below the 5th percentile (Barro, 2009; Backus et al., 2011; Beason and Schreindorfer, 2022). The results differ since I account for conditioning information embedded in the risk-neutral quantile function, whereas unconditional estimates of disaster risk tend to overestimate this risk, since the physical distribution acquires fatter tails when averaging out state variables.

Comparing the physical and risk-neutral quantile functions over time also sheds light on the role of risk aversion and forward-looking beliefs in jointly determining disaster risk premia. Particularly during crises, the value of an insurance contract that hedges against disaster risk increases, resulting in a downward movement in the left-tail of the risk-neutral quantile function. Simultaneously, investors often revise their beliefs about the likelihood of another disaster, frequently assigning a higher probability to such an event. This effect drives down the left-tail of the physical quantile function, creating an ambiguous overall impact on disaster risk premia. However, the quantile regression estimates indicate that the risk-neutral quantile function decreases proportionally more, highlighting the greater influence of risk aversion in determining disaster risk premia.

Given that the equity premium is primarily driven by disaster risk premia, the discussion above implies that the left-tail of the risk-neutral quantile function can predict the equity premium. An OLS regression of the equity premium against the 5% risk-neutral quantile shows preliminary evidence of forecasting ability, especially out-of-sample. In line with theoretical expectations, a decline in the 5% risk-neutral quantile is associated with a substantial increase in the equity premium. Notably, during the 2008 financial crisis and the 2020 Covid-19 crisis, monthly estimates of the equity premium reached peaks of around 5%.

Besides the equity premium puzzle, the QR estimates shed light on the role of first-order stochastic dominance and the pricing kernel puzzle. Specifically, I find that $\tilde{Q}_{t,\tau} < Q_{t,\tau}$ holds across most of the distribution, except in the far right-tail, where $\tilde{Q}_{t,\tau} > Q_{t,\tau}$ frequently occurs. This violation of stochastic dominance raises questions in asset pricing models using the expected utility framework, as it suggests that a representative investor exhibits negative risk aversion. Furthermore, I show that a violation of stochastic dominance implies that the pricing kernel is not monotonic, thereby confirming the pricing kernel puzzle while accounting for conditioning information, and without the need to estimate a density ratio.

To further understand the influence of disaster risk on the pricing kernel, I introduce a *distribution bound* on the SDF volatility that is closely related to the [Hansen and Jagannathan \(1991\)](#) bound. The distribution bound summarizes the risk-return trade-off of an asset paying out one dollar when the market return falls below a certain threshold. I show that disaster risk makes the risk-return trade off highly favorable by going short in an asset paying one dollar in case of a disaster. The price of such an asset is high because investors are willing to pay a significant premium

to insure against disaster risk, but the risk is limited since the actual probability of a disaster is comparatively low. The Sharpe ratio associated to this investment therefore dominates the Sharpe ratio of a direct investment in the market portfolio. Specifically, in the data, the Sharpe ratio on selling an asset that pays out one dollar if the return falls below the 5th percentile is 30% in monthly units, while the Sharpe ratio of investing in the market portfolio is only 13%. I also show that models which do not embed a source of disaster risk, such as conditional lognormal models, cannot rationalize this finding.

I conclude by proposing a model-free lower bound on disaster risk premia to assess the robustness of my earlier findings. This lower bound is observed from option prices and is inspired by recent bounds on the equity premium (Martin, 2017; Chabi-Yo and Loudis, 2020). Using quantile regression, I show that the lower bound explains a substantial proportion of the fluctuation in disaster risk premia over time. Moreover, the lower bound relaxes the assumption of a time-homogeneous relation between the physical and risk-neutral quantile functions. Empirically, the lower bound closely aligns with the disaster risk estimates derived from quantile regression, further strengthening the robustness of my earlier findings.

1.1 Related Literature

My approach, which uses quantile regression to estimate local dispersions between the physical and risk-neutral distribution, is related to a larger body of literature that estimates the pricing kernel from returns and option data (Ait-Sahalia and Lo, 2000; Jackwerth, 2000; Rosenberg and Engle, 2002; Beare and Schmidt, 2016; Linn et al., 2018; Cuesdeanu and Jackwerth, 2018). However, estimating the pricing kernel from returns and options can be challenging, especially in the tails of the distribution, where the ratio of densities that defines the pricing kernel can become unstable. In addition, using historical returns to estimate the physical density can lead to inconsistent results (Linn et al., 2018). Beason and Schreindorfer (2022) apply a similar methodology to decompose the unconditional equity premium.

In contrast, QR can be used to draw inference on the pricing kernel indirectly, by leveraging the *observed* realized return and risk-neutral distribution, which avoids the estimation of a density ratio. Furthermore, QR can account for changes in the shape and scale of the underlying SDF over time due to changing conditional information, while the approach of Cuesdeanu and Jackwerth (2018) renders an estimate of the SDF that only allows the normalizing constant to be time-varying, since the shape and scale are time invariant (see Section 4.1).

I use the QR estimates to shed light on disaster risk, by analyzing the Lorenz curve associated to the equity premium. Effectively, quantiles decompose the equity premium state-by-state, which is conceptually different from prior literature such as [Schneider \(2019\)](#) and [Chabi-Yo and Loudis \(2023\)](#), whose decompositions rely on averages across the return distribution. This state-by-state decomposition allows one to diagnose more precisely which part of the return distribution contributes to the equity premium. Unlike [Schneider \(2019\)](#), I also incorporate time series data, which enhances estimation of the physical measure. [Chabi-Yo and Loudis \(2023\)](#) is more closely aligned with my methodology since they decompose the market return into three distinct components. However, the estimation of the physical measure differs significantly, as I rely on quantile regression rather than non-linear weighted least squares. Notably, quantile regression proves to be more robust, especially in the tails. [Beason and Schreindorfer \(2022\)](#) also employ a state-by-state decomposition but provide unconditional estimates, potentially leading to an overestimation of disaster risk. Specifically, my methodology shows that the worst 5% of returns contribute only 17% to the equity premium, while [Beason and Schreindorfer \(2022\)](#) report an estimate of 91.5%.

My approach to infer conditional disaster risk is also related to the high-frequency literature. [Bollerslev and Todorov \(2011\)](#) and [Bollerslev et al. \(2015\)](#) use semi-martingale theory to dissect which part of the equity premium is coming from downside/upside risk. This approach is very general but limited to short horizons. Furthermore, it separates the contribution from diffusion and jump risk, while the QR approach can also be used for any part of the return distribution. The QR approach does not require the specification of underlying state variables driving the economy, as these variables are implicitly embedded within the risk-neutral quantile function. A similar rationale has been applied by [Andersen et al. \(2015\)](#) in a high-frequency context, leveraging derivative prices to gain insights into qualitative features of latent state variables.

Complementary to the QR estimates, I derive a nonparametric bound on the SDF volatility closely related to the bound of [Hansen and Jagannathan \(1991\)](#). They argue that the SDF is necessarily volatile and use this observation to screen asset pricing models. Several papers have built on this insight using higher-order moment bounds ([Snow, 1991](#); [Almeida and Garcia, 2012](#); [Liu, 2021](#)) and entropy bounds ([Stutzer, 1995](#); [Bansal and Lehmann, 1997](#); [Alvarez and Jermann, 2005](#); [Backus et al., 2014](#)). These bounds all provide a measure of how much the risk-neutral distribution differs from the physical distribution. Unlike the distribution bound, all of these measures

are global in that they rely on averages over the entire state space. The distribution bound in this paper is a function rather than a single statistic and can be considered an intermediate approach between a single bound and a complete estimate of the SDF.

The last part of this paper is also related to the growing literature on using options to estimate forward-looking equity premia (Martin, 2017; Martin and Wagner, 2019; Chabi-Yo and Loudis, 2020). However, unlike those papers that focus on the conditional expectation of excess returns, this paper uses option data to bound conditional return quantiles. Since the obtained bound does not require any parameter estimation and provides information on the entire distribution, it complements the recovery literature (Ross, 2015; Schneider and Trojani, 2019). Furthermore, the observed time variation in the approximation for the left-tail quantile, as documented in this paper, aligns with the concept of time-varying disaster risk proposed in various models by Gabaix (2012), Wachter (2013), Constantinides and Ghosh (2017), Isoré and Szczerbowicz (2017), Farhi and Gourio (2018) and Seo and Wachter (2019).

The rest of this paper is organized as follows. Section 2 presents the main empirical results from the quantile regressions and its consequences for the equity premium and SDF are discussed in Section 3. Section 4 provides further evidence on the robustness of QR to estimate disaster risk relative to extant approaches. Section 5 introduces the distribution bound, discusses its use in asset pricing models, and presents estimates of the distribution bound from empirical data. Building on the results of Sections 2 and 3, Section 6 establishes a model-free lower bound on disaster risk premia. Finally, Section 7 concludes.

2 Empirical Estimates of Quantile Difference

This section documents empirical estimates of the conditional difference between the physical and risk-neutral quantile functions. I first discuss the notation and then consider an example to clarify the idea and motivate the methodology.

2.1 Notation

Let $R_{m,t \rightarrow N}$ denote the market return from period t to $t + N$, where N typically represents 30-, 60-, or 90-days. The risk-free rate over the same period is denoted by $R_{f,t \rightarrow N}$, which is assumed to be known at time t . In the absence of arbitrage, there exists a positive random variable $M_{t \rightarrow N}$ such that, conditional on the investor's information at time t ,

$$\mathbb{E}_t [M_{t \rightarrow N} R_{m,t \rightarrow N}] = 1. \quad (2.1)$$

The random variable $M_{t \rightarrow N}$ is referred to as the stochastic discount factor (SDF) and the expectation in (2.1) is calculated under the *physical* probability measure \mathbb{P}_t , which is the actual distribution of the market return, i.e. $R_{m,t \rightarrow N} \sim \mathbb{P}_t$. The SDF can potentially depend on many state variables, but these are suppressed from the notation for brevity. It is convenient to restate (2.1) in terms of risk-neutral probabilities:

$$\tilde{\mathbb{E}}_t(R_{m,t \rightarrow N}) = 1/\mathbb{E}_t[M_{t \rightarrow N}] = R_{f,t \rightarrow N},$$

where the expectation is calculated under the *risk-neutral* measure $\tilde{\mathbb{P}}_t$ induced by $M_{t \rightarrow N}$. Finally, $F_t(x) := \mathbb{P}_t(R_{m,t \rightarrow N} \leq x)$ denotes the physical CDF of the market return conditional on the investor's information available at time t , $f_t(\cdot)$ denotes the conditional probability density function (PDF) and $Q_{t,\tau}$ denotes the conditional τ -quantile. As before, a tilde superscript refers to the risk-neutral measure, so that

$$\tilde{F}_t(\tilde{Q}_{t,\tau}) = \tilde{\mathbb{P}}_t(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau}) = \tau, \quad \forall \tau \in (0, 1).$$

The physical and risk-neutral quantiles depend on the underlying random variable $R_{m,t \rightarrow N}$ (i.e., $\tilde{Q}_{t,\tau} := \tilde{Q}_{t,\tau}(R_{m,t \rightarrow N})$), but I typically omit this dependence as the underlying random variable always refers to the market return.

To clarify my approach of using quantiles to analyze disaster risk, I consider the following asset pricing model that will be used several times in the paper.

Example 2.1 (Disaster risk). Consider the disaster risk model analyzed in [Backus et al. \(2011\)](#). The SDF process is given by

$$\log M_{t \rightarrow N} = \log \beta - \gamma \log G_{t \rightarrow N},$$

where β is a time discount factor, γ is the coefficient of relative risk aversion, and $G_{t \rightarrow N} = C_{t+N}/C_t$ is consumption growth in period $t + N$. Consumption growth follows a two-component structure:

$$\log G_{t \rightarrow N} = z_{1,t+N} + z_{2,t+N}, \quad z_{1,t+N} \sim \mathcal{N}(\mu, \sigma^2),$$

and $z_{2,t+N}$ is a Poisson mixture of normals to capture jumps representing rare shocks to consumption growth that are large in magnitude. The number of jumps, j , take on nonnegative integer values with probability $e^{-\omega} \omega^j / j!$, and conditional on j , the jump term is normal: $z_{2,t+N} | j \sim \mathcal{N}(j\theta, j\delta^2)$. [Backus et al. \(2011\)](#) show that the risk-neutral distribution of consumption growth in a representative agent model is again a normal mixture with parameters:

$$\tilde{\mu} = \mu - \gamma\sigma^2, \quad \tilde{\omega} = \omega e^{-\gamma\theta + (\gamma\delta)^2}, \quad \tilde{\theta} = \theta - \gamma\delta^2. \quad (2.2)$$

In this setup, risk aversion amplifies the jump frequency ($\tilde{\omega} > \omega$ if $\theta < 0$) as well as the jump size ($\tilde{\theta} < \theta$). If the model is calibrated such that $\theta \ll 0$, then $z_{2,t+N}$ can be interpreted as a disaster shock if a jump takes place ($j \geq 1$).

Figures 1a and 1b illustrate the impact of jumps on the physical and risk-neutral quantile functions. Specifically, in the absence of jumps, the market return follows a lognormal distribution and Figure 1a shows that the difference between the physical and risk-neutral quantile functions is approximately equal in both tails. However, when jumps are introduced, this difference is almost entirely concentrated in the left-tail. This result is driven by the impact of jumps on the risk-neutral distribution, and the requirement that $\theta < 0$ is crucial to drive a wedge between the physical and risk-neutral measures in the left-tail (see (2.2)). The question is whether these distinct shape restrictions on the physical and risk-neutral distribution are supported by the data.

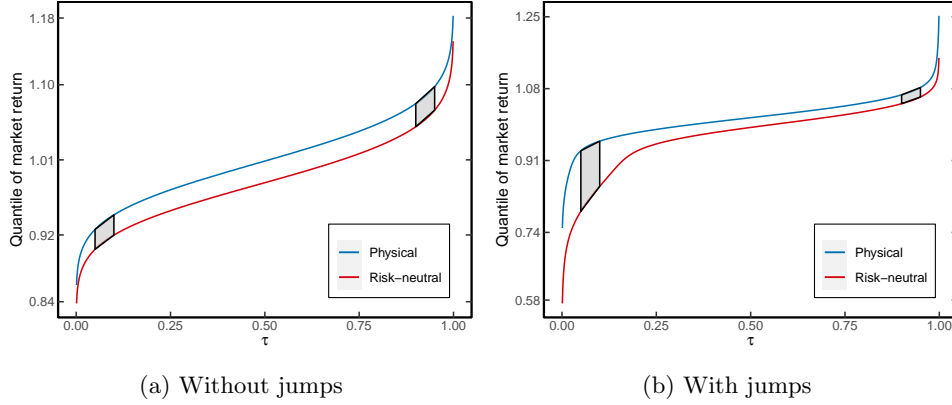


Figure 1: **Effect of jumps on physical and risk-neutral quantile functions.** The left panel displays the physical and risk-neutral quantile functions without jumps ($\omega = 0$), while the right panel illustrates the quantile functions with jumps ($\omega = 1.4$). In both cases, the mean of the disaster shock (θ) is set to -0.0074 . The market return is defined as a levered claim on the consumption asset. The trapezoids represent the difference in quantile functions at the 10th and 90th percentiles.

2.2 Methodology and Econometric Model

Building on the discussion in Example 2.1, it is of interest to estimate the quantile difference between the physical and risk-neutral measures. The disaster risk model predicts that these differences are significant in the left-tail while negligible in the right-tail. This is because investors' marginal utility of wealth in states associated

with a disaster is very high. Consequently, the following excess return,

$$\underbrace{\mathbb{E}_t [\mathbb{1}(R_{m,t \rightarrow N} \leq x)]}_{\text{price}} - \underbrace{\mathbb{E}_t [\mathbb{1}(R_{m,t \rightarrow N} \leq x)]}_{\text{expected payoff}} = \tilde{F}_t(x) - F_t(x),$$

can be rather high for thresholds x in the left-tail. The price is high because investors are willing to pay for insurance against disaster risk, even though the actual probability of a disaster can be significantly lower. This effect drives a wedge between the physical and risk-neutral CDFs in the left-tail. For analytical and estimation convenience, it proves more fruitful to consider the inverse of the CDFs (i.e., the quantile functions). Therefore, I consider $Q_{t,\tau} - \tilde{Q}_{t,\tau}$, and refer to these differences in the left-tail as *disaster risk premia*.

While the conditional risk-neutral distribution and its quantile function can be inferred from option prices without specific modeling assumptions (Breedon and Litzenberger, 1978), the same cannot be said for the physical distribution, unless strong assumptions are made about the martingale component of the SDF (Ross, 2015; Borovička et al., 2016). The information available about the conditional physical distribution is limited to a single realization of the market return, as $R_{m,t \rightarrow N}$ follows \mathbb{P}_t conditional on time t . Consequently, the primary challenge in measuring disaster risk premia lies in the unobservable nature of $Q_{t,\tau}$, which has made model-free inference challenging thus far.

2.2.1 Risk-Neutral Quantile Regression

In order to overcome this difficulty, I assume the following model for the physical quantile function

$$\underbrace{Q_{t,\tau}(R_{m,t \rightarrow N})}_{\text{Unobserved}} = \beta_0(\tau) + \beta_1(\tau) \underbrace{\tilde{Q}_{t,\tau}(R_{m,t \rightarrow N})}_{\text{Observed}}, \quad \forall \tau \in (0, 1). \quad (2.3)$$

If the world is risk-neutral, $[\beta_0(\tau), \beta_1(\tau)] = [0, 1]$ for all τ . Departures from risk-neutrality at a specific percentile τ are reflected by point estimates of $[\beta_0(\tau), \beta_1(\tau)]$ that are far from the $[0, 1]$ benchmark. Given a sample of T observations $\{R_{m,t \rightarrow N}, \tilde{Q}_{t,\tau}\}_{t=1}^T$, the unknown parameters in (2.3) can be estimated by quantile regression (Koenker and Bassett, 1978):

$$[\hat{\beta}_0(\tau), \hat{\beta}_1(\tau)] = \arg \min_{(\beta_0, \beta_1) \in \mathbb{R}^2} \sum_{t=1}^T \rho_\tau(R_{m,t \rightarrow N} - \beta_0 - \beta_1 \tilde{Q}_{t,\tau}), \quad (2.4)$$

where $\rho_\tau(\cdot)$ is the check function from quantile regression

$$\rho_\tau(x) = \begin{cases} \tau x, & \text{if } x \geq 0 \\ (\tau - 1)x & \text{if } x < 0. \end{cases}$$

Even if the world is not risk-neutral, the model in (2.3) can still be correctly specified, as is the case for conditional lognormal models (see Section 4.1). When the model is misspecified, the estimation in (2.4) remains meaningful as QR finds the best linear approximation to the conditional quantile function (Angrist et al., 2006).² Since the risk-neutral quantile itself is a highly non-linear transformation of state variables, the model can accommodate non-linear dependence between the physical quantile function and state variables driving the economy. The benefit of using the risk-neutral quantile function as a regressor is that it does not require the econometrician to take a stand on the state variables driving the physical distribution. Furthermore, both $R_{m,t \rightarrow N}$ and $\tilde{Q}_{t,\tau}$ are conditioned on the same information set, thus avoiding the mismatched information critique of Linn et al. (2018).

In addition, theory often suggests tantalizing links between the tails of the physical and risk-neutral distribution. Table 1 presents correlations between $Q_{t,\tau}$ and $\tilde{Q}_{t,\tau}$ for both left and right tails in different asset pricing models. In most models, these correlations are nearly one, indicating a strong positive relation that can be modeled by (2.3). Only for $\tau = 0.3$, the correlation is notably lower at 41% in the Campbell and Cochrane (1999) model and -67% in the Drechsler and Yaron (2011) model.

In Appendix B.1, I consider non-linear specifications as alternatives to the linear model in (2.3). Broadly speaking, I find that the linear model outperforms all non-linear models when predicting the physical quantile function out-of-sample. Based on this evidence, and the close linear approximation suggested by asset pricing models, I use the linear specification throughout most of the paper.

Remark 1. An alternative to QR is nonparametric estimation of the SDF as proposed by Aït-Sahalia and Lo (2000), Jackwerth (2000) and Rosenberg and Engle (2002). This method can infer the quantile difference from the estimated SDF but relies on pooled historical returns, which can be problematic for forward-looking distribution estimation (Linn et al., 2018). More recently, Linn et al. (2018) and Cuesdeanu

²This result is analogous to OLS, which finds the best linear approximation to the conditional expectation function, even if the model is misspecified.

Table 1: **Tail correlations (in %) of physical and risk-neutral quantile function in asset pricing models**

Percentile	0.05	0.1	0.2	0.3	0.7	0.8	0.9	0.95
Lognormal								
Campbell and Cochrane (1999)	96.94	94.49	83.61	40.88	86.51	94.25	97.27	98.26
Bansal and Yaron (2004)	99.97	99.97	99.98	99.98	99.99	99.99	99.99	100.00
Disaster								
Drechsler and Yaron (2011)	99.90	99.44	94.67	-67.16	96.88	98.75	99.45	99.67
Wachter (2013)	95.40	99.63	99.57	98.98	99.71	99.88	99.94	99.97
Constantinides and Ghosh (2017)	99.86	99.72	99.21	97.58	85.96	94.68	97.32	97.90

Note: This table reports the correlation between $Q_{t,\tau}$ and $\tilde{Q}_{t,\tau}$ in conditional lognormal models and models that embed a source of conditional disaster risk. The correlations at various percentiles are obtained by simulating 10^6 draws of the ergodic distribution of states in each model.

and [Jackwerth \(2018\)](#) proposed an estimator of the SDF that accounts for forward-looking information. However, this method presents challenges such as non-convex optimization, the objective function might be undefined due to the small number of existing risk-neutral moments (see Figure G8), ambiguity in basis function selection, and the inability to account for shape changes in the SDF leading to incorrect conditional inference. QR, on the other hand, avoids these issues, as shown in more detail in Section 4.1.

2.2.2 Measures of Fit

Based on the quantile regression (2.4), I consider two measures of fit to evaluate how well the risk-neutral quantile locally approximates the physical distribution. The first in-sample measure, $R^1(\tau)$, is defined as³

$$R^1(\tau) := 1 - \frac{\min_{b_0, b_1} \sum_{t=1}^T \rho_\tau(R_{m,t \rightarrow N} - b_0 - b_1 \tilde{Q}_{t,\tau})}{\min_{b_0} \sum_{t=1}^T \rho_\tau(R_{m,t \rightarrow N} - b_0)}. \quad (2.5)$$

This measure of fit was proposed by [Koenker and Machado \(1999\)](#) and is a clean substitute for the OLS R^2 . I also consider an out-of-sample measure of fit

$$R_{oos}^1(\tau) := 1 - \frac{\sum_{t=w}^T \rho_\tau(R_{m,t \rightarrow N} - \tilde{Q}_{t,\tau})}{\sum_{t=w}^T \rho_\tau(R_{m,t \rightarrow N} - \bar{Q}_{t,\tau})}, \quad (2.6)$$

where $\bar{Q}_{t,\tau}$ is the historical rolling quantile of the market return from time $t - w + 1$ to t , and w is the rolling window length. Notice that (2.6) is a genuine out-of-sample metric since no parameter estimation is used. In the equity premium literature, [Campbell and Thompson \(2008\)](#) stress the importance of out-of-sample predictability; (2.6) is analogous to their out-of-sample R^2 .

³It is well known that b_0 in the denominator of (2.5) equals the in-sample τ -quantile.

2.3 Data and Estimation

To estimate the quantile regression in (2.4), I require data on the market return and the risk-neutral distribution over time. I use overlapping returns on the S&P500 index from WRDS over the period 2003–2021 to represent the market return. I calculate the market return over a horizon of 30-, 60-, and 90-days. Second, over the same horizon, I use put and call option prices on the S&P500 on each day t from OptionMetrics to estimate the risk-neutral quantile function based on the [Breedon and Litzenberger \(1978\)](#) formula:

$$\tilde{F}_t\left(\frac{K}{S_t}\right) = R_{f,t \rightarrow N} \frac{\partial}{\partial K} \text{Put}_t(K), \quad (2.7)$$

where $\text{Put}_t(K)$ denotes the time t price of a European put option on the S&P500 index with stock price S_t , strike price K and expiration date $t + N$. This formula is model-free and only requires a no-arbitrage assumption. Due to the lack of a continuum of option prices, interpolation of different maturity options and missing data for option prices far in- and out-of-the money, it is a nontrivial exercise to obtain accurate estimates of \tilde{F}_t (and hence $\tilde{Q}_{t,\tau}$) from (2.7). A detailed description of my approach that overcomes these issues is described in Appendix C.2, which is based on [Filipović et al. \(2013\)](#).⁴ Finally, I obtain the risk-free rate from Kenneth French’s website.⁵

Table 2 shows the QR estimates of (2.4). The point estimates are close to the $[0, 1]$ benchmark in the right-tail ($\tau \geq 0.7$), but not in the left-tail ($\tau \leq 0.3$).⁶ Additionally, the joint restriction that $[\beta_0(\tau), \beta_1(\tau)] = [0, 1]$ is rejected for all $\tau \leq 0.2$, at all horizons. In contrast, the null hypothesis is never rejected for $\tau \geq 0.8$. The fact that the risk-neutral distribution provides a good approximation of the physical distribution in the right-tail is confirmed by the measures of fit, $R^1(\tau)$ and $R_{oos}^1(\tau)$, which are also shown in Table 2. Specifically, both in- and out-of-sample, the risk-neutral quantile fits the physical distribution much better in the right-tail.

Remark 2. The standard errors for the quantile regression in Table 2 are obtained by the smooth extended tapered block bootstrap (SETBB) of [Gregory et al. \(2018\)](#), which is robust to heteroscedasticity and weak dependence.⁷ This robustness is

⁴This approach uses a kernel density and adds several correction terms to approximate the risk-neutral density. I follow [Barletta and Santucci de Magistris \(2018\)](#) and use a principal components step to avoid overfitting in the tails.

⁵See http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html#Research

⁶Because the risk-neutral quantile function is estimated, there is a concern for attenuation bias due to measurement error. Unreported simulations show that this bias is very small in a setting that mimics the empirical application.

⁷It may seem counterintuitive that the standard errors decrease in the tails, which are generally

important in the estimation, since I use overlapping returns which creates time dependence in the error term, akin to the overlapping observation problem in OLS (Hansen and Hodrick, 1980). SETBB also renders an estimate of the covariance matrix between $\hat{\beta}_0(\tau)$ and $\hat{\beta}_1(\tau)$, which can be used to test joint restrictions on the coefficients.⁸

3 Equity Premium Puzzle and SDF Implications

Building on the estimates in Table 2, this section shows that the conditional equity premium is driven by disaster risk, and that disaster risk is a pervasive feature of the data, which poses a new challenge to asset pricing models. I further comment on two implications of Table 2 that relate to properties of the SDF that have previously received attention in the literature.

3.1 Equity Premium Puzzle

The results in Table 2 show that the physical distribution is close to risk-neutral in the right-tail, but not in the left-tail. Investors in the market portfolio thus get compensated for bearing downside risk, but not upside risk. This result has important repercussions for explanations of the equity premium puzzle. To see this, consider the following decomposition of the equity premium⁹

$$\begin{aligned} \mathbb{E}_t [R_{m,t \rightarrow N}] - R_{f,t \rightarrow N} &= \int_0^1 (Q_{t,\tau} - \tilde{Q}_{t,\tau}) d\tau \\ &= \underbrace{\int_0^{\underline{\tau}} (Q_{t,\tau} - \tilde{Q}_{t,\tau}) d\tau}_{\text{disaster risk}} + \int_{\underline{\tau}}^1 (Q_{t,\tau} - \tilde{Q}_{t,\tau}) d\tau, \end{aligned} \quad (3.1)$$

where $\underline{\tau}$ is a percentile close to zero. The first term on the right-hand side aggregates the local difference between the risk-neutral and physical quantiles in the left-tail, which I define as the contribution of disaster risk. The results in Table 2 show that these differences are the primary determinant for the equity premium, as in the right-tail we have $Q_{t,\tau} \approx \tilde{Q}_{t,\tau}$. The latter finding is consistent with the modeling assumption in (time-varying) disaster risk models that shocks to the market return

harder to estimate. However, since the regressor $\tilde{Q}_{t,\tau}$ changes with τ , there is an opposing effect that can cause the standard errors to decrease in the tails. This happens if $\tilde{Q}_{t,\tau}$ is more variable in the tails, akin to the intuition in OLS that more variability in the regressor decreases the standard error. In the data, $\tilde{Q}_{t,\tau}$ is much more variable in the tails.

⁸I use the `QregBB` function from the eponymous *R*-package, available on the author's Github page: <https://rdrr.io/github/gregorkb/QregBB/man/QregBB.html>. The only user required input for this method is the block length in the bootstrap procedure.

⁹See Appendix A.1 for a derivation.

Table 2: **Risk-neutral quantile regression**

Horizon	τ	$\hat{\beta}_0(\tau)$	$\hat{\beta}_1(\tau)$	Wald test (p -value)	$R^1(\tau)$ [%]	$R^1_{\text{OOS}}(\tau)$ [%]	Hit[%]	$\hat{Q}_{t,\tau} > \tilde{Q}_{t,\tau}$ [%]
30 days*	0.05	0.43 (0.208)	0.56 (0.223)	0.00	6.28	6.11	-2.67 (0.676)	99.88
	0.1	0.45 (0.201)	0.54 (0.209)	0.01	3.45	1.01	-3.56 (1.089)	98.52
	0.2	0.69 (0.284)	0.30 (0.290)	0.02	0.55	0.89	-3.73 (1.719)	90.98
	0.3	1.02 (0.357)	-0.02 (0.360)	0.00	0.00	2.49	-5.51 (2.147)	99.58
	0.4	1.17 (0.237)	-0.16 (0.237)	0.00	0.03	1.75	-7.32 (2.357)	97.25
	0.6	-0.45 (0.216)	1.44 (0.213)	0.00	4.62	4.19	-8.05 (2.468)	99.93
	0.7	-0.18 (0.162)	1.18 (0.159)	0.03	7.79	7.47	-5.84 (2.220)	99.95
	0.8	-0.09 (0.141)	1.09 (0.137)	0.19	12.44	12.50	-3.24 (1.886)	99.95
	0.9	0.03 (0.113)	0.97 (0.108)	0.96	20.41	21.88	-0.04 (1.235)	55.85
	0.95 (Obs. 4333)	0.12 (0.120)	0.89 (0.114)	0.57	27.07	31.31	0.27 (0.863)	22.41
60 days**	0.05	0.45 (0.303)	0.54 (0.343)	0.00	3.12	13.14	-3.33 (0.875)	100.00
	0.1	0.58 (0.263)	0.41 (0.283)	0.00	1.79	3.50	-5.57 (1.320)	100.00
	0.2	0.78 (0.336)	0.21 (0.345)	0.01	0.38	-0.03	-6.60 (2.351)	99.95
	0.3	0.93 (0.434)	0.07 (0.438)	0.00	0.01	-0.12	-7.81 (3.012)	99.47
	0.4	0.36 (0.325)	0.65 (0.323)	0.02	0.25	2.34	-8.48 (3.439)	99.79
	0.6	-0.65 (0.342)	1.64 (0.333)	0.02	5.57	4.60	-7.68 (3.465)	99.77
	0.7	-0.31 (0.266)	1.30 (0.256)	0.05	8.41	7.65	-7.34 (3.260)	99.91
	0.8	-0.08 (0.183)	1.08 (0.174)	0.07	12.70	12.23	-5.53 (2.683)	100.00
	0.9	0.04 (0.147)	0.96 (0.138)	0.58	21.66	22.79	-1.94 (1.707)	92.86
	0.95 (Obs. 4312)	0.04 (0.135)	0.96 (0.126)	0.90	31.07	34.19	0.43 (1.046)	13.73
90 days***	0.05	0.60 (0.405)	0.37 (0.478)	0.01	2.90	15.63	-2.95 (1.102)	100.00
	0.1	0.59 (0.321)	0.40 (0.356)	0.00	3.46	3.84	-6.36 (1.495)	100.00
	0.2	0.57 (0.516)	0.43 (0.534)	0.03	0.83	1.93	-7.53 (2.896)	100.00
	0.3	0.62 (0.637)	0.39 (0.643)	0.04	0.17	-0.52	-8.42 (3.668)	99.84
	0.4	0.42 (0.468)	0.60 (0.463)	0.02	0.22	-1.76	-9.52 (4.199)	99.77
	0.6	-0.84 (0.426)	1.82 (0.413)	0.01	6.37	3.81	-11.60 (4.542)	99.98
	0.7	-0.46 (0.307)	1.45 (0.293)	0.02	10.45	8.87	-9.43 (4.056)	100.00
	0.8	-0.23 (0.204)	1.23 (0.192)	0.10	15.47	16.54	-6.66 (3.189)	100.00
	0.9	-0.02 (0.170)	1.02 (0.157)	0.79	23.18	27.92	-1.12 (1.971)	100.00
	0.95 (Obs. 4291)	0.08 (0.153)	0.93 (0.139)	0.86	32.14	39.88	-0.06 (1.366)	52.37

Note: This table reports the QR estimates of (2.4) over the sample period 2003–2021 at different horizons, using overlapping returns. Standard errors are shown in parentheses and based on SETBB with a block length equal to the prediction horizon. *Wald test* denotes the p -value of the joint restriction $[\beta_0(\tau), \beta_1(\tau)] = [0, 1]$. $R^1(\tau)$ denotes the goodness of fit measure (2.5). $R^1_{\text{OOS}}(\tau)$ is the out-of-sample goodness of fit (2.6), using a rolling window of size 10 times the prediction horizon. Hit refers to the sample expectation defined in (3.3) and standard errors are reported in parentheses, which are obtained by stationary bootstrap based on 10,000 bootstrap samples. The last column indicates the time series average of the event that $\hat{Q}_{t,\tau} > \tilde{Q}_{t,\tau}$, where $\hat{Q}_{t,\tau} = \hat{\beta}_0(\tau) + \hat{\beta}_1(\tau)\tilde{Q}_{t,\tau}$.

are negative conditional on a disaster occurring (see, e.g., the condition $\theta < 0$ in Example 2.1). Hence, an asset pricing model seeking to explain the (conditional) equity premium of the market return must embed a source of disaster risk.

To illustrate the pervasiveness of disaster risk in the data, I consider the Lorenz

curve associated with the conditional equity premium

$$L_t(x) := \frac{\int_0^x (Q_{t,\tau} - \tilde{Q}_{t,\tau}) d\tau}{\mathbb{E}_t[R_{m,t \rightarrow N}] - R_{f,t \rightarrow N}} \stackrel{(3.1)}{=} \frac{\int_0^x (Q_{t,\tau} - \tilde{Q}_{t,\tau}) d\tau}{\int_0^1 (Q_{t,\tau} - \tilde{Q}_{t,\tau}) d\tau}, \quad 0 \leq x \leq 1.$$

The Lorenz curve summarizes the proportion of the equity premium contributed by the bottom $x\%$ of returns, akin to its interpretation in labor economics to summarize wealth inequality. Since $Q_{t,\tau}$ is unobserved, I use instead the inferred value, $\hat{Q}_{t,\tau} = \hat{\beta}_0(\tau) + \hat{\beta}_1(\tau)\tilde{Q}_{t,\tau}$, with the estimated parameters coming from the QR estimates in (2.4).

Figure 2a shows the average Lorenz curve in the data, together with the Lorenz curve implied by various asset pricing models.¹⁰ In the data, the Lorenz curve is quite concave, thus showing that the majority of the equity premium is contributed by the left-tail. At the same time, my estimation adds nuance to the degree of disaster risk influencing the equity premium. Specifically, while the disaster risk models of Barro (2009) and Backus et al. (2011) attribute approximately 90% of the equity premium to the lowest 5% of returns, empirical estimates suggest this proportion is only around 17%. These findings also deviate substantially from the nonparametric estimates of Beason and Schreindorfer (2022), who report that 91.5% of the equity premium is driven by the bottom 5% of returns. Our results differ because I account for conditioning information, while Beason and Schreindorfer (2022) employ an unconditional approach. Using unconditional averages can inflate the tails of the physical distribution (Chabi-Yo et al., 2008), leading to an overestimation of disaster risk.

On the other hand, the models of Campbell and Cochrane (1999) and Bansal and Yaron (2004) are even more misspecified since the Lorenz curve in these models is slightly convex, thus attributing more than 50% of the equity premium to upside returns. The model of Schreindorfer (2020) matches the Lorenz curve best, even though it also overestimates the contribution of disaster risk to the equity premium.

I also consider the Gini coefficient derived from the Lorenz curve

$$G_t = 2 \int_0^1 L_t(\tau) d\tau - 1.$$

By construction, the Gini coefficient is between -1 and 1, and a value closer to 1 indicates that a bigger proportion of the equity premium is coming from the left-tail.

¹⁰I thank Beason and Schreindorfer (2022) for making the code to simulate from these models publicly available.

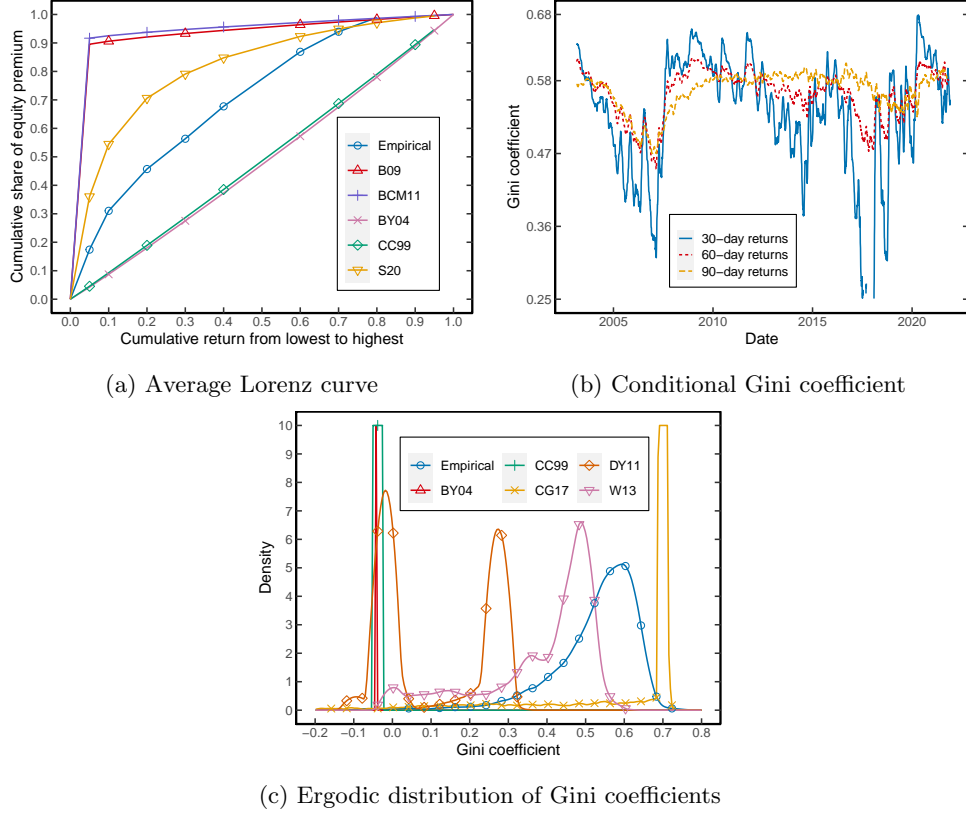


Figure 2: **Lorenz curve and Gini coefficient of the conditional equity premium.** This figure presents the Lorenz curve and Gini coefficient associated with the conditional equity premium in both empirical data and asset pricing models. Panel (a) displays the time-averaged Lorenz curve estimated from 30-day returns (Empirical) alongside Lorenz curves implied by the unconditional asset pricing models of Barro (2009) (B09), Backus et al. (2011) (BCM11), and Schreindorfer (2020) (S20), as well as the average Lorenz curve from conditional asset pricing models by Campbell and Cochrane (1999) (CC99) and Bansal and Yaron (2004) (BY04). Panel (b) depicts the estimated Gini coefficient over time for different return horizons and is smoothed using a 30-day rolling window. Panel (c) shows the ergodic distribution of Gini coefficients estimated from 30-day returns (Empirical) and those implied by the conditional asset pricing models of CC99, BY04, Drechsler and Yaron (2011) (DY11), Wachter (2013) (W13), and Constantinides and Ghosh (2017) (CG17). Model parameters are calibrated on a monthly frequency, and the ergodic distribution is derived from 10,000 state draws.

In contrast, a value of 0 suggests that the equity premium is evenly distributed across the return distribution, while negative values imply that the right-tails contribute more to the equity premium than the left-tails. Figure 2b shows the time series of conditional Gini coefficients for the various return horizons. For 30-day returns, the Gini coefficient mostly hovers between 0.33 and 0.68. At longer horizons, the Gini coefficients exhibit less variability and typically range between 0.47 to 0.6. These coefficients are also countercyclical, peaking during periods associated with economic downturns, such as the 2008 financial crisis and the Covid-19 crisis. Overall, the Gini coefficients consistently exhibit strong positive values, highlighting the pervasiveness

of conditional disaster risk in the data, which extends beyond crisis periods.

Finally, I analyze the ergodic distribution of Gini coefficients in time-varying asset pricing models and compare it to the distribution implied by the data.¹¹ Figure 2c displays these distributions and shows that many asset pricing models have difficulty in matching the empirical distribution. The conditional lognormal models of Campbell and Cochrane (1999) and Bansal and Yaron (2004) imply negative Gini coefficients, with minimal variation among different states, contrary to what the data indicate. The models of Drechsler and Yaron (2011), Wachter (2013), Constantinides and Ghosh (2017) all incorporate a source of disaster risk, but they also have difficulty to match the empirics. In particular, the models of Drechsler and Yaron (2011) and Wachter (2013) embed too little disaster risk, while the model of Constantinides and Ghosh (2017) overestimates the impact of disaster risk.

3.2 Driver of Disaster Risk Premia: Insurance or Beliefs?

Disaster risk premia have two components: an insurance effect and a forward-looking beliefs effect (under rational expectations). To see this, consider again the short position in a derivative security that pays one dollar if the market return is below a threshold, denoted by x , in the left-tail. The excess return on such an investment can be interpreted as

$$\underbrace{\mathbb{E}_t [\mathbb{1} (R_{m,t \rightarrow N} \leq x)]}_{\text{price of insurance}} - \underbrace{\mathbb{E}_t [\mathbb{1} (R_{m,t \rightarrow N} \leq x)]}_{\text{forward looking belief}} = \tilde{F}_t(x) - F_t(x).$$

During a crisis, the price of this insurance security tends to rise. This effect can occur in the disaster risk model (Example 2.1), if risk aversion increases when a disaster hits, leading to an increase in $\tilde{F}_t(x)$ and a subsequent decrease in $\tilde{Q}_{t,\tau}$. Simultaneously, investors may believe that the actual probability of a disaster increases during a crisis. This belief drives up $F_t(x)$ and, consequently, pushes down $Q_{t,\tau}$.

Building on this discussion, it is not immediately clear what the net effect is on disaster risk premia ($Q_{t,\tau} - \tilde{Q}_{t,\tau}$), as both $Q_{t,\tau}$ and $\tilde{Q}_{t,\tau}$ tend to decrease during periods of heightened market uncertainty. Figure 3a illustrates this effect for 30-day returns and $\tau = 0.05$. Notably, during the global financial crisis and Covid-19 crisis, both the physical and risk-neutral quantile functions exhibit significant drops.

To shed light on the net effect on disaster risk premia during crises, Figure 3b

¹¹In the model, I obtain the distribution of Gini coefficients from the state distribution. In the data, I rely on the time series average. If the data are generated by the model and the system is ergodic, Birkhoff's theorem implies that the state and time averages are equal almost everywhere.

displays the evolution of disaster risk premia over time. The most significant change occurs during the peak of the global financial crisis and the Covid-19 crisis. In these turbulent periods, disaster risk premia consistently rise, suggesting that the insurance effect is more substantial than the forward-looking beliefs effect.¹² Because of these large increases, disaster risk is a more important driver of the equity premium, which clarifies the countercyclical Gini coefficients in Figure 2b.

The downward fluctuations in the risk-neutral quantile function can be particularly pronounced, plummeting to as low as 63% during crisis periods. In contrast, the physical quantile function only drops to 78%, suggesting that a monthly loss of 22% or more had a 5% probability. To put this in perspective, this probability is 14 times higher than the estimate obtained from historical monthly S&P500 returns (from 1926 to 2021). This calculation shows that historical estimates can diverge significantly from forward-looking beliefs. Furthermore, the time fluctuations in the physical quantile function lend empirical support to the notion of time-varying disaster risk, as proposed in various models such as Gabaix (2012), Wachter (2013), Constantinides and Ghosh (2017), Isoré and Szczerbowicz (2017), Farhi and Gourio (2018) and Seo and Wachter (2019).

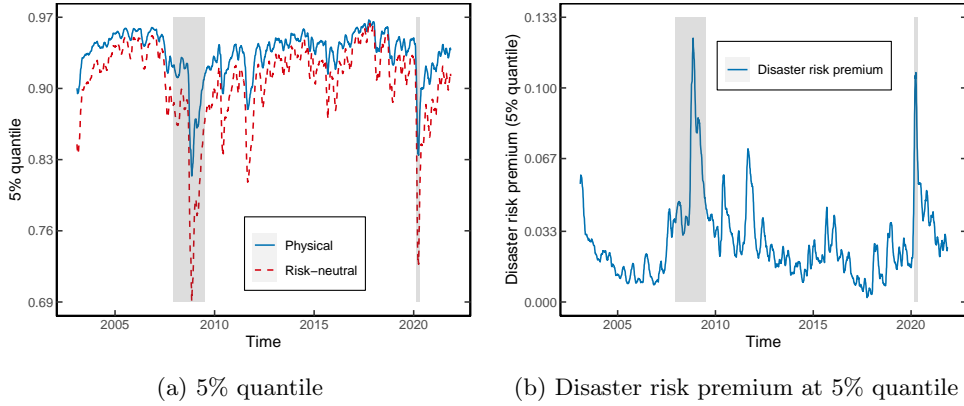


Figure 3: Disaster risk premia for 30-day returns at the 5th percentile. Panel (a) shows the physical and risk-neutral quantile functions over time at $\tau = 0.05$. The physical quantile function is estimated from the quantile regression in (2.4). Panel (b) shows the associated disaster risk premium, $Q_{t,\tau} - \tilde{Q}_{t,\tau}$. Both panels are smoothed using a 30-day moving window. The two shaded bars denote the Great Recession period (Dec 2007 – June 2009) and Covid-19 crisis (Feb 2020 – April 2020).

¹²I find similar results for 60- and 90-day returns.

3.3 Predicting the Equity Premium

The previous results establish that, in times of heightened market uncertainty, the equity premium is driven more by disaster risk. This observation suggests a strong link between $\mathbb{E}_t[R_{m,t \rightarrow N}] - R_{f,t \rightarrow N}$ and the tail of the risk-neutral distribution, which motivates the predictive regression

$$R_{m,t \rightarrow N} - R_{f,t \rightarrow N} = \beta_0 + \beta_1 \tilde{Q}_{t,\tau} + \varepsilon_{t \rightarrow T}, \quad (3.2)$$

where $\tilde{Q}_{t,\tau}$ is evaluated at $\tau = 0.05$.

Table 3 shows the results at several return horizons. In all cases, the coefficient is negative, consistent with previous findings that the equity premium increases under market uncertainty. Following Welch and Goyal (2008), the table also reports the out-of-sample R^2 , denoted by R_{oos}^2 , which compares the predictions of (3.2) to a rolling average of excess returns. Precisely, I estimate (3.2) using the sub-sample covering 2003–2012, and fix the estimated parameters to predict excess returns over the out-of-sample period 2013–2021. Encouragingly, R_{oos}^2 is always positive and statistically significant according to the Diebold and Mariano (1995) test, thus suggesting that the left-tail of the risk-neutral quantile function outperforms the historical mean benchmark. These values are also substantially higher compared to the R_{oos}^2 reported by Welch and Goyal (2008) using various valuation ratios, or Martin (2017) using SVIX.¹³

Figure 4 shows the estimated equity premium over time for 30- and 60-day returns. The panels are annualized to make them comparable. Both panels display considerable variation in the equity premium over time and large values during the global financial crisis and Covid-19 crisis. In these periods, Figure 4a suggests that the annualized equity premium peaks at 58%, which is substantial relative to more conventional estimates based on dividend-price ratios. On the other hand, the estimates around the 2008 financial crisis are in line with Martin (2017, Figure IV).

3.4 Pricing Kernel Monotonicity and Stochastic Dominance

Besides the equity premium puzzle, the QR estimates in Table 2 also provide insights into other asset pricing anomalies, such as pricing kernel monotonicity. Pricing kernel monotonicity refers to the property that $M_{t \rightarrow N}(R_{m,t \rightarrow N}) := \mathbb{E}[M_{t \rightarrow N} | R_{m,t \rightarrow N}]$ is a decreasing function of the market return. Asset pricing models that link the SDF to the marginal rate of substitution imply that the pricing kernel is indeed a

¹³The latter is not directly comparable however, since SVIX does not require parameter estimation.

Table 3: **OLS estimates of conditional equity premium**

Horizon	Full sample				Sub-sample		
	$\hat{\beta}_0$	$\hat{\beta}_1$	$R^2[\%]$	Obs	$R^2[\%]$	$R^2_{oos}[\%]$	p -value DM
30 days	0.13 (0.089)	-0.14 (0.096)	1.79	4333	10.44	1.82	0.00
60 days	0.17 (0.120)	-0.18 (0.137)	2.74	4312	18.16	3.30	0.00
90 days	0.20 (0.150)	-0.22 (0.178)	3.58	4291	26.67	4.20	0.00
Period	2003–2021				2013–2021		

Note: This table reports the OLS estimates of (3.2) for 30-, 60- and 90-day returns. Standard errors are shown in parentheses and calculated using stationary bootstrap, with an average block length equal to the return horizon. R^2_{oos} denotes the out-of-sample R^2 using the historical rolling mean of excess returns. The window length is equal to 5 years. p -value DM denotes the p -value of the Diebold and Mariano (1995) test that the risk-neutral quantile exhibits equal out-of-sample forecasting accuracy as the rolling mean. The “Period” row indicates the specific time periods used for estimation.




Figure 4: **Estimated equity premium.** This figure shows the estimated equity premium based on (3.2) for 30-day returns (Panel 4a) and 60-day returns (Panel 4b). In both cases, the equity premium is converted to annual units. The two shaded bars signify the Great Recession period (Dec 2007 – June 2009) and Covid-19 crisis (Feb 2020 – April 2020).

decreasing function. Empirically, there is suggestive evidence that the pricing kernel is not monotonic, which is puzzling as it contradicts that a representative investor is risk-averse (see Ait-Sahalia and Lo (1998), Jackwerth (2000), Rosenberg and Engle (2002), Bakshi et al. (2010), Beare and Schmidt (2016) and Cuesdeanu and Jackwerth (2018)). However, a formal statistical test that can detect violations of monotonicity is challenging as one needs uniform confidence bands for the estimated SDF, which requires tools from empirical process theory (see, e.g., Beare and Schmidt (2016)).

I consider a different approach based on stochastic dominance. Proposition A.1 in the Appendix shows that pricing kernel monotonicity implies that the physical distri-

bution is first-order stochastic dominant (FOSD) over the risk-neutral distribution, i.e., $F_t(x) \leq \tilde{F}_t(x)$ for all x . The latter condition can be rephrased as $F_t(\tilde{Q}_{t,\tau}) \leq \tau$ for all $\tau \in (0, 1)$. A violation of stochastic dominance, and hence pricing kernel monotonicity, is thus implied if there is statistical evidence that $F_t(\tilde{Q}_{t,\tau}) > \tau$ for a single τ . To investigate this possibility, let¹⁴

$$\begin{aligned}\text{Hit}_{t \rightarrow N} &= \mathbb{1} \left(R_{m,t \rightarrow N} < \tilde{Q}_{t,\tau} \right) - \tau, \\ \overline{\text{Hit}} &= \frac{1}{T} \sum_{t=1}^T \text{Hit}_{t \rightarrow N}.\end{aligned}\tag{3.3}$$

Hence, $\overline{\text{Hit}}$ provides an estimate of $\mathbb{E}(F_t(\tilde{Q}_{t,\tau}) - \tau)$ which ought to be negative for all τ under FOSD.¹⁵ The “Hit” column in Table 2 reports the value of (3.3), which is positive for $\tau = 0.95$ at the 30- and 60-day horizon. However, these estimates are not significant at the conventional levels and a violation of FOSD cannot be concluded.

Since $F_t(x) \leq \tilde{F}_t(x)$ if and only if $Q_{t,\tau} > \tilde{Q}_{t,\tau}$, it follows that violations of stochastic dominance can also be identified directly from the quantile function. Based on the QR estimates (2.4), consider the predicted quantile function $\hat{Q}_{t,\tau} = \hat{\beta}_0(\tau) + \hat{\beta}_1(\tau)\tilde{Q}_{t,\tau}$. The last column in Table 2 displays the time series average of instances where $\hat{Q}_{t,\tau} > \tilde{Q}_{t,\tau}$. Broadly speaking, for all horizons, violations of stochastic dominance are infrequent, except far in the right-tail. At $\tau = 0.95$, stochastic dominance is frequently violated, consistent with a non-monotonic pricing kernel.¹⁶ In representative agent models, this result is puzzling as it contradicts the assumption of decreasing marginal utility of wealth (see Proposition A.2 in the Appendix).

3.5 Belief Recovery

A recent literature asks to what extent Arrow prices can be used to learn about the underlying probability distribution of the data, or the subjective probabilities used by investors. Since Arrow prices are confounded by risk aversion, it is impossible to identify the underlying probabilities from Arrow prices alone, unless one imposes additional restrictions (Ross, 2015; Borovička et al., 2016; Bakshi et al., 2018; Qin et al., 2018; Jackwerth and Menner, 2020). For example, Ross (2015) uses the Perron-Frobenius theorem to recover investors’ beliefs, which agrees with the underlying physical measure under rational expectations.

¹⁴The $\text{Hit}_{t \rightarrow N}$ function was first introduced by Engle and Manganelli (2004) in a different context.

¹⁵ $\overline{\text{Hit}}$ also yields another measure of the difference between F_t and \tilde{F}_t . Consistent with the quantile regression estimates, the Hit statistic shows that F_t and \tilde{F}_t are similar in the right-tail, but different in the left-tail.

¹⁶The most significant violations occur during two major financial crises: the 2008 financial crisis and the 2020 Covid-19 crisis.

Complementary to this insight, the QR estimates in Table 2 show that the right-tail of the physical distribution can approximately be recovered from the right-tail of the risk-neutral distribution, which aligns with the investor’s belief under rational expectations. In contrast, the left-tail of the physical distribution cannot be recovered even though the risk-neutral quantile serves as a conservative lower bound. In Section 6, I propose a more stringent lower bound to recover the left-tail of the physical distribution as well from option data.

4 QR and Robust Estimation of Disaster Risk

Section 3.1 demonstrated that the conditional lognormal assumption is inconsistent with the observed disaster risk premia in the market. At the same time, Figure 2a showed that disaster risk models tend to overestimate the magnitude of disaster risk in the data. These conclusions heavily rely on the accuracy of QR in providing estimates of the physical quantile function.

In this section, I compare QR to nonparametric SDF methods for estimating disaster risk. Foreshadowing the results, I show that QR is more robust and argue that the SDF approach tends to overestimate disaster risk. These results help explain the current disagreement about the extent of disaster risk in the data, and provide further support for QR to estimate this risk.

4.1 QR in the Conditional Lognormal Model

To convey the intuition, it is convenient to work with a discretized version of the Black and Scholes (1973) model. There is a riskless asset that offers a certain return, $R_{f,t \rightarrow N} \equiv R_f = e^{r_f N}$, and a risky asset with return

$$R_{m,t \rightarrow N} = \exp([\mu_t - \frac{1}{2}\sigma_t^2]N + \sigma_t\sqrt{N}Z_{t+N}), \quad (4.1)$$

where μ_t represents the conditional mean return, σ_t is the conditional volatility, and Z_{t+N} is a random shock that follows a standard normal distribution. In this setup, $M_{t \rightarrow N} := \exp(-[r_f + \xi_t^2/2]N - \xi_t\sqrt{N}Z_{t+N})$ is a valid SDF with conditional Sharpe ratio

$$\xi_t = \frac{\mu_t - r_f}{\sigma_t}. \quad (4.2)$$

Hence, under risk-neutral measure, the conditional distribution of $R_{m,t \rightarrow N}$ is given by

$$\log \tilde{R}_{m,t \rightarrow N} \sim \mathcal{N}\left((r_f - \frac{1}{2}\sigma_t^2)N, \sigma_t^2 N\right). \quad (4.3)$$

Notice that σ_t is implicitly observed from the risk-neutral distribution, but μ_t is unobserved with mean $\mu := \mathbb{E}[\mu_t]$ and variance $\sigma_\mu^2 := \text{Var}(\mu_t) < \infty$. The following result characterizes the limiting behavior of the QR estimates (2.4) in the lognormal model when the variance of the equity premium is small. A convenient way to model this is by means of a drifting sequence $\sigma_\mu^T \rightarrow 0$ as $T \rightarrow \infty$, which captures the intuition that the volatility of the equity premium is much smaller than the return volatility.

Proposition 4.1 (QR in Lognormal Model). *In the lognormal model described above with return observations $\{R_{m,t \rightarrow N}\}_{t=1}^T$ and risk-neutral quantile functions $\{\tilde{Q}_{t,\tau}\}_{t=1}^T$, the following hold.*

- (i) *Suppose that conditional on time t , μ_t follows a normal distribution $\mu_t \sim \mathcal{N}(\mu, \sigma_\mu^2)$, independent of σ_t . Let $Q_{t,\tau}(\sigma_t, \sigma_\mu)$ denote the physical quantile function of $R_{m,t \rightarrow N}$ conditional on σ_t only. Then, for all $\tau \in \mathcal{I} :=$ a closed subset of $[\varepsilon, 1 - \varepsilon]$ for $0 < \varepsilon < 1$, the physical quantile function satisfies*

$$\begin{aligned} Q_{t,\tau}(\sigma_t, \sigma_\mu) &= \exp \left[\left(\mu - \frac{1}{2} \sigma_t^2 \right) N + \left(\sqrt{\sigma_\mu^2 N^2 + \sigma_t^2 N} \right) \Phi^{-1}(\tau) \right] \\ &= \tilde{Q}_{t,\tau} e^{(\mu - r_f)N} (1 + \mathcal{O}(\sigma_\mu N)), \end{aligned}$$

where $\Phi^{-1}(\tau)$ denotes the quantile function of the standard normal distribution.

- (ii) *Consider a drifting sequence for σ_μ , denoted by $\sigma_\mu^T \rightarrow 0$ as $T \rightarrow \infty$. Then, under Assumption A.4 in the Appendix, the estimated parameters in the quantile regression*

$$\left[\hat{\beta}_0(\sigma_\mu^T; \tau), \hat{\beta}_1(\sigma_\mu^T; \tau) \right] = \arg \min_{(\beta_0, \beta_1) \in \mathbb{R}^2} \sum_{t=1}^T \rho_\tau(R_{m,t \rightarrow N} - \beta_0 - \beta_1 \tilde{Q}_{t,\tau}),$$

satisfy

$$\left[\hat{\beta}_0(\sigma_\mu^T; \tau), \hat{\beta}_1(\sigma_\mu^T; \tau) \right] = \left[0, e^{(\mu - r_f)N} \right] + o_p(1). \quad (4.4)$$

Furthermore, the quantile forecast based on the QR estimates satisfies

$$\hat{\beta}_0(\sigma_\mu^T; \tau) + \hat{\beta}_1(\sigma_\mu^T; \tau) \tilde{Q}_{t,\tau} = Q_{t,\tau} + o_p(1). \quad (4.5)$$

Proof. See Appendix A.3. ■

Proposition 4.1(i) shows that the risk-neutral quantile function is a good predictor of $Q_{t,\tau}(\sigma_t; \sigma_\mu)$ when σ_μ is small, and the difference between the two functions is governed by the unconditional equity premium $e^{(\mu - r_f)N}$. In this case, Proposition 4.1(ii) suggests that the QR estimates are almost constant across τ and close

to $[0, e^{(\mu - r_f)N}]$. This result obtains without assuming that μ_t follows a normal distribution. The wedge between $Q_{t,\tau}(\sigma_t; \sigma_\mu)$ and $\tilde{Q}_{t,\tau}$ not explained by the equity premium can be attributed to uncertainty about μ_t , which increases the variance of the physical distribution. The assumption that σ_μ is small relative to σ_t accords with empirical findings of [Martin \(2017, Table I\)](#), who finds that $2.4\% \leq \sigma_\mu \leq 4.6\%$, whereas σ_t hovers around 20%. Unreported simulations show that the approximation in (4.4) obtains closely when the model is calibrated to match these stylized facts. As a result, the physical quantile forecast based on the QR estimates in (4.5) is also highly accurate.

4.2 QR versus Nonparametric SDF Estimation

Because of the availability of closed-form expressions in the lognormal model, it is instructive to compare the QR approach to alternative methods for estimating the physical distribution. Since the SDF represents the Radon–Nikodym derivative of the risk-neutral and physical measures, it is possible to obtain the physical quantile function from the estimated SDF. There is a substantial literature on how to estimate the SDF in a forward-looking manner (see [Remark 1](#)). For this comparison, I consider the state-of-the-art SDF estimator proposed by [Cuesdeanu and Jackwerth \(2018\)](#) (CJ).

After some algebra, the SDF in the Black-Scholes model can be expressed as a function of the market return:

$$M_{t \rightarrow N} = \exp \left(-\frac{N}{2} \left[\mu_t + r_f + \frac{r_f^2 - \mu_t^2}{\sigma_t^2} \right] \right) (R_{m,t \rightarrow N})^{-\xi_t / \sigma_t}, \quad (4.6)$$

where ξ_t is the conditional Sharpe ratio (4.2). CJ project the unobserved SDF in (4.6) on the market return and estimate an SDF of the form

$$\widehat{M}_{t \rightarrow N} = C_t g(R_{m,t \rightarrow N}),$$

where C_t is a time-varying constant, and $g(\cdot)$ is an unknown function that can be estimated by choosing a sieve basis. Since $g(\cdot)$ is *time-homogeneous*, it is evident that changes in the shape of the true SDF in (4.6) are not captured by the estimated SDF. Specifically, in times when the Sharpe ratio is high, the physical and risk-neutral measures exhibit more distinct differences, as the true SDF becomes steeper. Because the estimated SDF does not account for these shape changes, it leads to a severe underestimation of the physical quantile function in the left-tail. [Proposition 4.1](#) demonstrates that the QR approach does not suffer from this limitation.

To illustrate this discussion, I simulate returns from the lognormal model and estimate the physical quantile function at the 5th percentile using QR and the SDF estimate of CJ. Since the conditional (physical) quantile function is known analytically in the lognormal model, I evaluate the forecast accuracy using the quantile error ratio, $\hat{Q}_{t,\tau}/Q_{t,\tau}$, where $\hat{Q}_{t,\tau}$ is the predicted physical quantile based on QR or the SDF estimate. Panel 5a displays the empirical density of error ratios obtained by simulating 1,000 returns. In line with Proposition 4.1(ii), the error ratio corresponding to QR is symmetric and closely centered around one. In contrast, when the physical quantile is inferred from the estimated SDF, the error density is biased and exhibits fat tails since the estimated SDF cannot change shape. Consequently, in periods of high disaster risk premia, the CJ method severely *underestimates* $Q_{t,\tau}$.

Panel 5b presents the histogram of error ratios conditioned on the 30 largest values of $Q_{t,\tau} - \tilde{Q}_{t,\tau}$, clearly illustrating the downward bias in the SDF method. On average, the predicted physical quantile is 7% lower than its actual value when disaster risk premia are high. The QR approach is less affected by this bias because it can capture changes in the shape of the SDF. The computational benefits of QR are also notable, as the computation of the physical quantile forecast takes less than a second. On the other hand, the SDF method requires more than 20 minutes to complete the same task.¹⁷

The bottom panels of Figure 5 further illustrate the difference between QR and CJ using the 30-day return data from Section 2.3, particularly during the 2008 financial crisis and the Covid-19 crisis. At the height of both crises, both methods predict increases in disaster risk premia as $\hat{Q}_{t,\tau} - \tilde{Q}_{t,\tau}$ rises significantly. As mentioned earlier, the SDF approach implies that disaster risk premia increase less relative to the QR approach, as the shape of the SDF remains constant over time. However, it is worth noting that while the QR approach performs well when returns are conditional lognormal, Appendix B.2 demonstrates that the quantile forecasts based on QR contradict (4.5), casting further doubt on the validity of the conditional lognormal assumption in the data.

¹⁷Moreover, the optimization problem required to implement the sieve estimation did not converge, as the maximum number of iterations were exceeded. This problem occurs due to the large number of parameters to estimate, and because the optimization problem is not convex (see Remark 1).




Figure 5: **Disaster risk premia at the 5th percentile.** Panel (a) shows the quantile error ratio, $\hat{Q}_{t,\tau}/Q_{t,\tau}$, in a conditional [Black and Scholes \(1973\)](#) model for $\tau = 0.05$, where $\hat{Q}_{t,\tau}$ is the predicted physical quantile based on QR or the SDF estimate of CJ. Volatility is generated according to an AR(1)-model with mean value 0.2, standard deviation 0.03 and a persistency of 0.9. The mean of the physical distribution follows $\mu_t \sim \mathcal{N}(0.07, 0.02^2)$, the risk-free rate equals $r_f = 0.01$, the time horizon is one-year, and the number of observations is 1,000. Panel (b) shows the histogram of error ratios conditioned the 30 events for which $Q_{t,\tau} - \hat{Q}_{t,\tau}$ is maximal. The bottom panels illustrate the difference between the predicted physical quantile obtained from QR, and the predicted quantile coming from the SDF estimate of CJ, during the global financial crisis and the Covid-19 crisis. Both estimates are based on 30-day returns, using the data from [Section 2.3](#). The bottom panels are smoothed using a 30-day rolling window.

5 Disaster Risk and SDF Volatility

[Section 3.1](#) demonstrated that the physical and risk-neutral distributions locally differ most in the left-tail. In this section, I show that these local differences imply that the SDF must be highly volatile; an observation that is closely related to the [Hansen and Jagannathan \(1991\)](#) bound. Furthermore, I use this insight to argue that the left-tail of the physical distribution cannot be too predictable, which clarifies the low explanatory power in [Table 2](#).

5.1 A Bound on the SDF Volatility

For ease of notation, I define $\phi_t(\tau) := F_t(\tilde{Q}_{t,\tau})$, which can be interpreted as the ordinal dominance curve of the measures \mathbb{P}_t and $\tilde{\mathbb{P}}_t$ (Hsieh and Turnbull, 1996). Furthermore, let

$$\aleph_t^+ := \{M_{t \rightarrow N} : M_{t \rightarrow N} \geq 0 \text{ and } \mathbb{E}_t[M_{t \rightarrow N} R_{m,t \rightarrow N}] = 1\},$$

which is the space of all nonnegative conditional SDFs. The volatility bound on the SDF can now be stated as follows.

Proposition 5.1 (Distribution bound). *Assume no-arbitrage, then for any $M_{t \rightarrow N} \in \aleph_t^+$, we have*

$$\frac{\sigma_t(M_{t \rightarrow N})}{\mathbb{E}_t[M_{t \rightarrow N}]} \geq \frac{|\tau - \phi_t(\tau)|}{\sqrt{\phi_t(\tau)(1 - \phi_t(\tau))}} \quad \forall \tau \in (0, 1). \quad (5.1)$$

If a risk-free asset exists, then $\mathbb{E}_t[M_{t \rightarrow N}] = 1/R_{f,t \rightarrow N}$ and (5.1) simplifies to

$$\sigma_t(M_{t \rightarrow N}) \geq \frac{1}{R_{f,t \rightarrow N}} \frac{|\tau - \phi_t(\tau)|}{\sqrt{\phi_t(\tau)(1 - \phi_t(\tau))}} \quad \forall \tau \in (0, 1).$$

The bound can be further rewritten in terms of the conditional CDFs only

$$\sigma_t(M_{t \rightarrow N}) = \frac{1}{R_{f,t \rightarrow N}} \frac{|\tilde{F}_t(x) - F_t(x)|}{F_t(x)(1 - F_t(x))} \quad \forall x \in (0, \infty). \quad (5.2)$$

Proof. See Appendix A.4. ■

If $\mathbb{P}_t = \tilde{\mathbb{P}}_t$, agents are risk-neutral and the dominance curve evaluates to $\phi_t(\tau) = \tau$. In that case the distribution bound degenerates to zero. Proposition 5.1 makes precise the sense in which any local difference between the physical and risk-neutral distribution leads to a volatile SDF. Compare this to the classical Hansen and Jagannathan (1991) (HJ) bound:

$$\sigma_t(M_{t \rightarrow N}) \geq \frac{1}{R_{f,t \rightarrow N}} \frac{|\mathbb{E}_t[R_{m,t \rightarrow N}] - R_{f,t \rightarrow N}|}{\sigma_t(R_{m,t \rightarrow N})}. \quad (5.3)$$

The lower bound in (5.3) shows that any excess return leads to a volatile SDF. Essentially, (5.3) uses three sources of information: (i) the mean of the physical distribution (ii) the mean of the risk-neutral distribution (iii) the variance of the physical distribution. The lower bound in (5.3) is also a global measure of distance between \mathbb{P}_t and $\tilde{\mathbb{P}}_t$, since the mean and volatility are averages across the whole distribution.

In contrast, the bound in (5.2) compares the physical and risk-neutral distribution at every point x , which is a *local* measure of distance between \mathbb{P}_t and $\tilde{\mathbb{P}}_t$. To clarify

this local interpretation, consider the following decomposition of the (scaled) equity premium

$$\begin{aligned} \frac{\mathbb{E}_t[R_{m,t \rightarrow N}] - R_{f,t \rightarrow N}}{R_{f,t \rightarrow N}} &= -\mathbb{COV}_t[R_{m,t \rightarrow N}, M_{t \rightarrow N}] \\ &= \int_0^\infty \mathbb{COV}_t[\mathbb{1}(R_{m,t \rightarrow N} \leq x), M_{t \rightarrow N}] dx, \end{aligned} \quad (5.4)$$

where the first equation follows since the SDF prices the market return (2.1), and the second equation is a consequence of Hoeffding's identity (see Lemma A.5.). Equation (5.4) shows that $\mathbb{COV}_t[\mathbb{1}(R_{m,t \rightarrow N} \leq x), M_{t \rightarrow N}]$ locally measures the dependence between the SDF and market return. In other words, it quantifies how the SDF's variability relates to the market return's variability at different quantiles.

To explain the equity premium and disaster risk premia, the SDF must exhibit sufficient variability. Since the distribution bound can be derived from applying the Cauchy-Schwarz inequality to $\mathbb{COV}_t[\mathbb{1}(R_{m,t \rightarrow N} \leq x), M_{t \rightarrow N}]$, it is expected to yield sharper bounds on the SDF volatility than the HJ bound if, for example, there is high tail dependence between the SDF and market return such as in the disaster risk model.¹⁸

5.2 Quantile Predictability in the Left-Tail

The bound presented in Proposition 5.1 sheds light on the seemingly “low” explanatory power observed in the left-tail quantile regressions in Table 2. For tractability, it is more convenient to show this for CDFs instead of quantile functions, but the intuition remains the same. Specifically, suppose one could predict $F_t(x)$ at some x in the left-tail, then this prediction can be exploited by going short in an asset that pays $\mathbb{1}(R_{m,t \rightarrow N} \leq x)$. The profit and risk associated to this investment are, respectively

$$\begin{aligned} &\frac{1}{R_{f,t \rightarrow N}} \left(\tilde{\mathbb{E}}_t[\mathbb{1}(R_{m,t \rightarrow N} \leq x)] - \mathbb{E}_t[\mathbb{1}(R_{m,t \rightarrow N} \leq x)] \right) \\ &= \frac{1}{R_{f,t \rightarrow N}} \left(\tilde{F}_t(x) - F_t(x) \right), \end{aligned} \quad (5.5)$$

$$\sigma_t(\mathbb{1}(R_{m,t \rightarrow N} \leq x)) = \sqrt{F_t(x)(1 - F_t(x))}.$$

Although such binary state payoffs do not exist in reality, they can be replicated closely by a portfolio of put options. In consequence, high predictability of $F_t(x)$ in the left-tail would render too good a Sharpe ratio; a near-arbitrage opportunity.

¹⁸See McNeil et al. (2015, Chapter 7.2.4) for a formal definition of tail dependence.

Following the reasoning in [Ross \(2005, Chapter 5\)](#), a crude upper bound on the SDF volatility imposes limitations on the degree of predictability in the left-tail by the distribution bound in [Proposition 5.1](#). This argument breaks down in the right-tail since [\(5.5\)](#) is roughly zero, and high predictability would not imply counterfactually high SDF volatility.

5.3 Distribution Bound in Asset Pricing Models

The estimated Gini coefficients in [Section 3.1](#) demonstrate that conditional disaster risk is a pervasive feature of the data. This section complements those findings using the unconditional version of the distribution bound in [Proposition 5.1](#):

$$\frac{\sigma(M)}{\mathbb{E}[M]} \geq \frac{\tau - \phi(\tau)}{\sqrt{\phi(\tau)(1 - \phi(\tau))}}, \quad (5.6)$$

where $\sigma(M)$ represents the unconditional SDF volatility, and $\phi(\tau) = F(\tilde{Q}_\tau)$. In this context, $F(\cdot)$ denotes the unconditional physical CDF, and \tilde{Q}_τ is the unconditional risk-neutral quantile function of the market return. The main benefit of using unconditional distributions is that they can be estimated without running the risk-neutral quantile regressions. Moreover, the bound in [\(5.6\)](#) only requires the estimation of distribution functions, whereas existing approaches typically use unconditional density functions to estimate disaster risk (see, e.g. [Beason and Schreindorfer \(2022\)](#)).

The subsequent examples demonstrate that the HJ bound is always stronger than the distribution bound in models that do not embed a source of disaster risk. In contrast, models that incorporate disaster risk can generate distribution bounds that exceed the HJ bound in the left-tail. Since I use unconditional distributions, the time subscripts will be omitted from the notation.

Example 5.1 (CAPM). The Capital Asset Pricing Model (CAPM) specifies the SDF as

$$M = \alpha - \beta R_m,$$

where R_m denotes the return on the market portfolio. In this case $M \notin \mathbb{N}^+$, since the SDF can become negative. However, this probability is very small over short time horizons or we can think of M as an approximation to $M^* := \max(0, M) \in \mathbb{N}^+$. Since the HJ bound is derived by applying the Cauchy-Schwarz inequality to $\mathbb{C}\mathbb{O}\mathbb{V}(R_m, M)$, the inequality binds if M is a linear combination of R_m . Hence, under CAPM, the HJ bound is (weakly) stronger than the distribution bound regardless of the distribution of R_m .

Example 5.2 (Joint normality). Suppose that M and R_m are jointly normally distributed and denote the mean and variance of R_m by μ_R and σ_R^2 respectively.

The normality assumption violates no-arbitrage since M can be negative, but could be defended as an approximation over short time horizons when the variance is small (see Example 5.3). In Appendix A.5, I prove that

$$\left| \text{COV} \left(\mathbb{1} \left(R_m \leq \tilde{Q}_\tau \right), M \right) \right| = f_R(\tilde{Q}_\tau) |\text{COV}(R_m, M)|, \quad (5.7)$$

where $f_R(\cdot)$ is the marginal density of R_m .¹⁹ This identity gives an explicit expression for the weighting factor in Hoeffding's identity (5.4). In Appendix A.5, I also derive an explicit expression for the relative efficiency between the distribution and HJ bound, defined by

$$\frac{\text{HJ bound}}{\text{distribution bound}} = \frac{\sqrt{\phi(\tau)(1-\phi(\tau))}}{\sigma_R f_R(\tilde{Q}_\tau)}. \quad (5.8)$$

To see that the HJ bound is always stronger than the distribution bound, minimize (5.8) with respect to τ . Appendix A.6 shows that the minimizer τ^* satisfies $\tilde{Q}_{\tau^*} = \mu_R$. For this choice, $\phi(\tau^*) = \mathbb{P}(R_m \leq \tilde{Q}_{\tau^*}) = 1/2$ and $f_R(\tilde{Q}_{\tau^*}) = 1/\sqrt{2\pi\sigma_R^2}$. Therefore, (5.8) can be bounded by

$$\frac{\sqrt{\phi(\tau)(1-\phi(\tau))}}{\sigma_R f(\tilde{Q}_\tau)} \geq \frac{\sqrt{2\pi}}{2} \approx 1.25.$$

Hence, the HJ bound is always stronger in a model where the SDF and market return are jointly normal.

Example 5.3 (Joint lognormality). Let Z_R and Z_M be standard normal random variables with correlation ρ and consider the specification

$$\begin{aligned} R_m &= e^{(\mu_R - \frac{\sigma_R^2}{2})N + \sigma_R \sqrt{N} Z_R} \\ M &= e^{-(r_f + \frac{\sigma_M^2}{2})N + \sigma_M \sqrt{N} Z_M}, \end{aligned}$$

where N governs the time scale in annual units. Simple algebra shows that the no-arbitrage condition, $\mathbb{E}[MR_m] = 1$, is satisfied when $\mu_R - r_f = -\rho\sigma_R\sigma_M$. It is difficult to find an analytical solution for the relative efficiency between the HJ and distribution bound in this case, but linearization leads to a closed form expression which is quite accurate in simulations. The details are described in Appendix A.7, where I show that

$$\min_{\tau \in (0,1)} \frac{\text{HJ bound}}{\text{distribution bound}} \approx \frac{1}{2} \sqrt{\frac{2\pi\sigma_R^2 N}{\exp(\sigma_R^2 N) - 1}}. \quad (5.9)$$

This expression is independent of μ_R . An application of l'Hôpital's rule reveals that

¹⁹Notice that this is the marginal density under physical measure \mathbb{P} .

the relative efficiency converges to $\sqrt{2\pi}/2$ if $N \rightarrow 0^+$.²⁰ The ratio in (5.9) is less than 1 if $\sigma_R \geq 0.92$ and $N = 1$. Since the annualized market return volatility is about 16%, the HJ bound is stronger than the distribution bound under any reasonable parameterization if the SDF and market return are lognormal.

Example 2.1 (Continued). The disaster risk model discussed in Section 2.1 is calibrated according to the results in Backus et al. (2011, Table II). The market return in this model is considered as a levered claim on consumption growth, i.e. an asset that pays dividends proportional to $G_{t \rightarrow N}^\lambda$. Here λ governs the variability of the claim to equity. I convert the model implied volatility bounds to monthly units, to facilitate the comparison with the empirical bounds obtained in Section 5.5.

The distribution bound, HJ bound and SDF volatility are depicted in Panels 6a (without jumps) and 6b (with jumps). Consistent with Example 5.3, the distribution bound in the model without jumps never exceeds the HJ bound because both the market return and SDF follow lognormal distributions. The distribution bound with jumps has a sharp peak at $\tau = 0.037$, after which it steadily decreases. Interestingly, there is a range of τ values for which the distribution bound is stronger than the HJ bound.²¹ This result can be understood from the physical and risk-neutral quantile functions in Figure 1b. The risk-neutral quantile function displays a heavy left-tail, owing to the implied disaster risk embedded in the SDF. Consequently, it is extremely profitable to sell digital put options which pay out in case of a disaster. These put options must have high Sharpe ratios as their prices are high (insurance against disaster risk), but the actual probability of a disaster event occurring is low enough that the risk associated with selling such insurance is limited.

5.4 Data and Empirical Estimation of the Distribution Bound

To further illustrate the presence of disaster risk in the data, I estimate the distribution bound (5.1) empirically, using the same 30-day S&P500 returns as discussed in Section 2.3. However, in this case, I use non-overlapping returns that cover the period 1996–2021.²² These returns are sampled at the middle of each month, resulting in a total of 312 observations. Over this period, the Sharpe ratio is 13%, and the HJ bound therefore implies that the monthly SDF is quite volatile.

²⁰This is the same relative efficiency in Example 5.2, which is unsurprising as the linearization becomes exact in the limit as $N \rightarrow 0^+$.

²¹In Appendix A.8, I show that the distribution bound can also exceed the HJ bound when returns follow the Pareto distribution.

²²I use non-overlapping returns in this section to facilitate testing and to make the results comparable to other nonparametric bounds, which are typically estimated based on non-overlapping returns (see e.g. Liu (2021)).

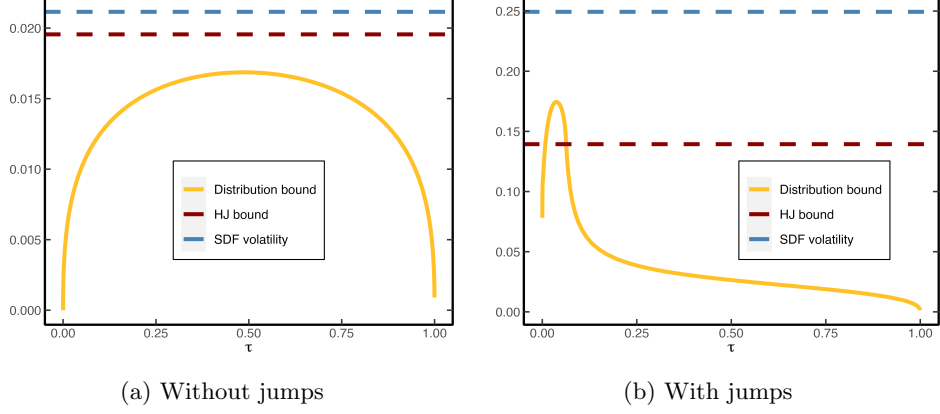


Figure 6: **HJ and distribution bound in disaster risk model without and with jumps.** Panels (a) and (b) show the HJ and distribution bound for the disaster risk model (Example 2.1) without and with jumps, respectively. The bounds and true SDF volatility are reported in monthly units. Parameters are calibrated according to Backus et al. (2011, Table II).

The distribution bound consists of three unknowns that need to be estimated: (i) the physical distribution (F); (ii) the risk-neutral quantile function (\tilde{Q}_τ), and; (iii) the risk-free rate (R_f). To estimate the unconditional risk-free rate, denoted by \hat{R}_f , I rely on the historical average of monthly interest rates. Next, to obtain an estimate of the physical distribution, I employ a kernel (CDF) estimator, given by:

$$\hat{F}(x) := \frac{1}{T} \sum_{t=1}^T \Phi \left(\frac{x - R_{m,t \rightarrow N}}{h} \right), \quad (5.10)$$

where $\Phi(\cdot)$ is the Epanechnikov kernel and h is the bandwidth determined by cross-validation. This choice of estimator ensures that the distribution bound is a smooth function of τ , which reduces the impact of outliers relative to the discontinuous empirical CDF.

Finally, I apply the procedure outlined in Section 2.3 to estimate \tilde{F}_t (the conditional risk-neutral CDF). Subsequently, I average the conditional distributions to estimate the unconditional CDF:

$$\hat{\tilde{F}}(x) := \frac{1}{T} \sum_{t=1}^T \tilde{F}_t(x).$$

Under appropriate assumptions about the distribution of returns, $\hat{\tilde{F}}$ converges to \tilde{F} as $T \rightarrow \infty$. An estimate of the unconditional risk-neutral quantile function can then be obtained from

$$\hat{\tilde{Q}}(\tau) := \inf \left\{ x \in \mathbb{R} : \tau \leq \hat{\tilde{F}}(x) \right\}. \quad (5.11)$$

Finally, based on the physical CDF (5.10) and risk-neutral quantile function (5.11), I estimate the distribution bound by

$$\hat{\theta}(\tau) := \frac{|\tau - \hat{\phi}(\tau)|}{\sqrt{\hat{\phi}(\tau)(1 - \hat{\phi}(\tau))\hat{R}_f}}, \quad \tau \in [\varepsilon, 1 - \varepsilon] \subseteq (0, 1), \quad (5.12)$$

where $\hat{\phi}(\tau) := \hat{F}(\hat{Q}(\tau))$ is the estimated ordinal dominance curve and ε is a small positive number.

5.5 Unconditional Evidence of Disaster Risk

Figure 7a illustrates the estimated physical and risk-neutral measures, which differ most in the left-tail. The distribution bound shows that this difference leads to a volatile SDF, which is shown in Figure 7b. The lower bound on the SDF volatility implied by the distribution bound is much stronger than the HJ bound in the left-tail. This finding aligns with empirical evidence documenting that high Sharpe ratios can be attained by selling out-of-the money put options (see Broadie et al. (2009) and the references therein). The supremum of the distribution bound occurs around the 5th percentile, implying that the monthly SDF volatility must exceed 31%. This value is more than twice the level indicated by the sample HJ bound. Moreover, the shape of the distribution bound is quite similar to the distribution bound implied by the disaster risk model in Figure 6b.²³




Figure 7: **Physical/risk-neutral CDF and distribution bound for monthly S&P500 returns.** Panel (a) shows the unconditional physical and risk-neutral CDF for monthly S&P500 returns, over the period 1996-2021. Panel (b) shows the distribution bound as function of τ , together with the HJ bound.

²³The non-monotonicity in the right-tail of the distribution bound occurs because $\tilde{F}(x) > F(x)$, for x large enough. That is, the physical distribution does not first-order stochastically dominates the risk-neutral distribution. This result is consistent with the positive $\hat{\text{Hit}}$ estimates in Table 2.

The graphical evidence suggests that the distribution bound renders a stronger bound on the SDF volatility than the HJ bound. To test this hypothesis more formally, I fix a priori the probability level at 0.037 ($\tau = 0.037$), which renders the sharpest bound on the SDF volatility in the disaster risk model (Example 2.1). At this probability level, the distribution bound is 26% in the data, which is roughly double the level implied by the HJ bound.

To see whether this difference is statistically significant, I consider the following test statistic

$$\mathcal{T} := \hat{\theta}(0.037) - \frac{|\bar{R}_m - \hat{R}_f|}{\hat{\sigma}\hat{R}_f}. \quad (5.13)$$

The first term on the right denotes the estimated distribution bound (5.12) evaluated at the 3.7th percentile, using the entire time series of returns $\{R_{m,t \rightarrow N}\}$. The second term denotes the estimated HJ bound, using \bar{R}_m and $\hat{\sigma}$ as the respective sample mean and standard deviation of $\{R_{m,t \rightarrow N}\}$. A value of $\mathcal{T} > 0$ indicates that the distribution bound is stronger than the HJ bound. To test this restriction, consider the null and alternative hypothesis:

$$\begin{aligned} H_0 : \mathcal{T} &\leq 0 \\ H_1 : \mathcal{T} &> 0. \end{aligned} \quad (5.14)$$

Since the distribution of (5.13) is difficult to characterize, I use stationary bootstrap to approximate the p -value under the null hypothesis. The stationary bootstrap is used to generate time indices from which we recreate (with replacement) bootstrapped returns $\{R_{m,t \rightarrow N}^*\}$ (Politis and Romano, 1994). The same bootstrapped time indices are used to re-estimate the physical CDF and risk-neutral quantile function. I repeat the bootstrap exercise 100,000 times and for each bootstrap sample, I calculate the test statistic \mathcal{T}^* . Finally, the empirical p -value is obtained as the fraction of times $\mathcal{T}^* \leq 0$. The last column in Table 4 shows that the p -value is 7.5%, which provides preliminary evidence that the distribution bound significantly exceeds the HJ bound in the left-tail.

Table 4: **Sample bounds and bootstrap result**

Sample size	HJ bound	distribution bound	p -value
312	0.133	0.260	0.075

Note: This table reports the HJ and distribution bound for monthly S&P500 returns over the period 1996–2021. The distribution bound is evaluated at $\tau = 0.0374$. The final column denotes the p -value of the null hypothesis in (5.14). The p -value is obtained from 100,000 bootstrap samples and counts the fraction of times that $\mathcal{T}^* \leq 0$.

Remark 3. When the HJ bound is stronger than the distribution bound, many of the bootstrap samples may not include disaster shocks. Over the entire sample period, there are only two instances where returns were less than -20%: in September 2008 and February 2020. When considering bootstrap samples that include both of these months, the p -value is only 3.6%. In contrast, the p -value increases to 22% for bootstrap samples that exclude these months. These findings underscore the sensitivity of the test to the presence of disaster shocks. Overall, the results suggest that, unconditionally, the SDF needs to be highly volatile to be consistent with local differences between the physical and risk-neutral measure in the left-tail.

6 A Model-Free Lower Bound on Disaster Risk Premiums

The previous findings indicate that the risk-neutral quantile function is not a good approximation of the physical quantile function in the left-tail. In this section, I derive a lower bound on disaster risk premiums observed from option prices. This lower bound does not require parameter estimation and relaxes the assumption of a time-homogeneous linear relation between the physical and risk-neutral quantiles in (2.3).

6.1 Approximating the Quantile Difference

To analyze the difference between $Q_{t,\tau}$ and $\tilde{Q}_{t,\tau}$, I use some elementary tools from functional analysis. The quantile function can be regarded as a map φ between normed spaces, taking as input a distribution function and returning the quantile function: $\varphi(F_t) = F_t^{-1} = Q_{t,\tau}$. Expanding φ around the observed risk-neutral CDF yields

$$Q_{t,\tau} - \tilde{Q}_{t,\tau} = \varphi(F_t) - \varphi(\tilde{F}_t) = \varphi'_{\tilde{F}_t}(F_t - \tilde{F}_t) + o\left(\|F_t - \tilde{F}_t\|\right), \quad (6.1)$$

where $\|\cdot\|$ is a norm on a suitable linear space²⁴ and $\varphi'_{\tilde{F}_t}(F_t - \tilde{F}_t)$ is the Gâteaux derivative of φ at \tilde{F}_t in the direction of F_t :

$$\begin{aligned} \varphi'_{\tilde{F}_t}(F_t - \tilde{F}_t) &:= \lim_{\lambda \downarrow 0} \frac{\varphi\left[(1-\lambda)\tilde{F}_t + \lambda F_t\right]}{\lambda} \\ &= \frac{\partial}{\partial \lambda} \varphi\left((1-\lambda)\tilde{F}_t + \lambda F_t\right) \Big|_{\lambda=0}. \end{aligned} \quad (6.2)$$

Heuristically, the Gâteaux derivative can be thought of as measuring the change

²⁴Formally, the space can be defined as $\{\Delta : \Delta = c(F - G), F, G \in \mathbb{D}, c \in \mathbb{R}\}$ and \mathbb{D} is the space of distribution functions (Serfling, 2009). See van der Vaart (2000, Section 20.1) and Serfling (2009, p. 217) for further details about the approximation.

in the quantile function when the risk-neutral distribution is moved in the direction of the physical distribution. Appendix A.9 shows that the Gâteaux derivative is given by

$$\varphi'_{\tilde{F}_t}(F_t - \tilde{F}_t) = \frac{\tau - F_t(\tilde{Q}_{t,\tau})}{\tilde{f}_t(\tilde{Q}_{t,\tau})} = \frac{\tau - \phi_t(\tau)}{\tilde{f}_t(\tilde{Q}_{t,\tau})}, \quad (6.3)$$

where $\phi_t(\tau) = F_t(\tilde{Q}_{t,\tau})$ denotes the conditional ordinal dominance curve. I proceed under the working hypothesis that the remainder term in (6.1) is “small” in the sup-norm, $\|g\|_\infty = \sup_x |g(x)|$.

Assumption 6.1. *The remainder term in (6.1) can be neglected.*

Remark 4. The assumption implies that the first order approximation in (6.1) is accurate. The condition that $\|F_t - \tilde{F}_t\|_\infty$ is small can be understood as excluding near-arbitrage opportunities, since the distribution bound in Proposition 5.1 shows that substantial pointwise differences between $F_t(\cdot)$ and $\tilde{F}_t(\cdot)$ lead to a very volatile SDF.

I combine (6.1) and (6.3) in conjunction with Assumption 6.1 to obtain the approximation

$$Q_{t,\tau} - \tilde{Q}_{t,\tau} \approx \underbrace{\frac{\tau - F_t(\tilde{Q}_{t,\tau})}{\tilde{f}_t(\tilde{Q}_{t,\tau})}}_{\text{risk-adjustment}}. \quad (6.4)$$

The second term on the right can be thought of as a risk-adjustment term to capture the unobserved wedge between $Q_{t,\tau}$ and $\tilde{Q}_{t,\tau}$. The approximation in (6.4) contains the terms $\tilde{Q}_{t,\tau}$ and $\tilde{f}_t(\tilde{Q}_{t,\tau})$, which are directly observed at time t using the Breeden and Litzenberger (1978) formula in (2.7). However, $F_t(\cdot)$ is unknown and hence (6.4) cannot be used directly to approximate $Q_{t,\tau}$.

6.2 A Lower Bound on Disaster Risk Premia

To make further progress, I show that the numerator term, $\tau - F_t(\tilde{Q}_{t,\tau})$, can be lower bounded with option data under economically motivated constraints. This bound, combined with the approximation in (6.4), will then imply a lower bound on disaster risk premia.

I start from the observation that the SDF in representative agent models can be expressed as a function of the market return (Chabi-Yo and Loudis, 2020):

$$\frac{\mathbb{E}_t[M_{t \rightarrow N}]}{M_{t \rightarrow N}} = \frac{\frac{u'(W_t x_0)}{u'(W_t x)}}{\tilde{\mathbb{E}}_t\left[\frac{u'(W_t x_0)}{u'(W_t x)}\right]} \quad \text{with } x = R_{m,t \rightarrow N} \text{ and } x_0 = R_{f,t \rightarrow N}, \quad (6.5)$$

where W_t is the agent's wealth at time t and $u(x)$ represents the agent's utility function. Define

$$\zeta(x) := \frac{u'(W_t R_{f,t \rightarrow N})}{u'(W_t x)} \quad \text{and} \quad \theta_k = \frac{1}{k!} \left(\frac{\partial^k \zeta(x)}{\partial x^k} \right)_{x=R_{f,t \rightarrow N}}. \quad (6.6)$$

Notice that $\zeta(\cdot)$ is simply the inverse of the intertemporal marginal rate of substitution (IMRS) and θ_k are the coefficients of its Taylor expansion around $R_{f,t \rightarrow N}$. I make the following assumptions about the market return and the IMRS of the representative agent.

Assumption 6.2. *In the representative agent model, it holds that (i) $\tilde{\mathbb{E}}_t [R_{m,t \rightarrow N}^3] < \infty$; and (ii) $\zeta^{(4)}(x) \leq 0$.*

Assumption 6.2(i) allows for fat tails in the risk-neutral distribution as long as the third moment exists. This assumption relaxes the implicit assumption made by Chabi-Yo and Loudis (2020) that infinitely many moments exist. Figure G8 in the Appendix illustrates that the risk-neutral distribution frequently exhibits a finite number of moments, some of which may not exceed 4, particularly in turbulent market conditions. Chabi-Yo and Loudis (2020) present sufficient conditions for 6.2(ii) to hold, which relate to the sign of the fifth derivative of the utility function of the representative agent. Specifically, for common utility functions such as CRRA or HARA utility, parameter restrictions are needed to ensure that 6.2(ii) holds.²⁵

I need one additional assumption to bound disaster risk premia. To state this assumption and the resulting lower bound, I use the following notation for high-order risk-neutral moments and truncated high-order risk-neutral moments, respectively.

$$\begin{aligned} \tilde{\mathbb{M}}_{t \rightarrow N}^{(n)} &:= \tilde{\mathbb{E}}_t [(R_{m,t \rightarrow N} - R_{f,t \rightarrow N})^n] \\ \tilde{\mathbb{M}}_{t \rightarrow N}^{(n)}[k_0] &:= \tilde{\mathbb{E}}_t [\mathbb{1}(R_{m,t \rightarrow N} \leq k_0) (R_{m,t \rightarrow N} - R_{f,t \rightarrow N})^n]. \end{aligned} \quad (6.7)$$

Assumption 6.3. *In the representative agent model, the following holds:*

- (i) $(-1)^{k-1} \theta_k \geq \frac{1}{R_{f,t \rightarrow N}^k}$ for $k = 1, 2, 3$
- (ii) $\tilde{\mathbb{M}}_{t \rightarrow N}^{(3)} \leq 0$.

Chabi-Yo and Loudis (2020, Table 6) provide empirical evidence that 6.3(i) holds with equality when estimating the conditional equity premium. Assumption 6.3(ii) is a very mild restriction on risk-neutral skewness, which is almost always negative

²⁵For example, for CRRA utility, the risk aversion coefficient cannot be too large. See Appendix D for a detailed discussion.

at every date and time horizon. This empirical fact is well known.²⁶

The following two propositions show how option data can be employed to establish bounds on the difference between the physical and risk-neutral measures in the left-tail.

Proposition 6.4 (Lower Bound on CDF). *Suppose Assumptions 6.2 and 6.3 hold, and assume that the risk-neutral density exists. Then,*

$$\tau - F_t(\tilde{Q}_{t,\tau}) \geq \frac{\sum_{k=1}^3 \frac{(-1)^{k-1}}{R_{f,t \rightarrow N}^k} \left(\tau \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)} - \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)}[\tilde{Q}_{t,\tau}] \right)}{1 + \sum_{k=1}^3 \frac{(-1)^{k-1}}{R_{f,t \rightarrow N}^k} \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)}} =: \text{CLB}_{t,\tau}, \quad (6.8)$$

for all $\tau \leq \tau'$, where τ' is defined implicitly by

$$\tilde{Q}_{t,\tau'} = \min \left(R_{f,t \rightarrow N} - \sqrt{\widehat{\text{VAR}}_t(R_{m,t \rightarrow N})}, \tilde{Q}_{t,\tau^*} \right),$$

and \tilde{Q}_{t,τ^*} is defined in Theorem A.12.

Proof. See Appendix A.10. ■

Proposition 6.5 (Lower Bound on Disaster Risk Premia). *Consider the same assumptions in Proposition 6.4 and assume additionally that Assumption 6.1 holds. Then, for all $\tau \leq \tau'$*

$$Q_{t,\tau} - \tilde{Q}_{t,\tau} \geq \overbrace{\frac{\text{CLB}_{t,\tau}}{f_t(\tilde{Q}_{t,\tau})}}^{\text{risk-adjustment}} =: \text{LB}_{t,\tau}. \quad (6.9)$$

Proof. By Assumption 6.1, the approximation in (6.4) holds, which in combination with Proposition 6.4 renders

$$\begin{aligned} Q_{t,\tau} - \tilde{Q}_{t,\tau} &\stackrel{(6.4)}{\approx} \frac{\tau - F_t(\tilde{Q}_{t,\tau})}{f_t(\tilde{Q}_{t,\tau})} \\ &\stackrel{(6.8)}{\geq} \frac{1}{f_t(\tilde{Q}_{t,\tau})} \left(\frac{\sum_{k=1}^3 \frac{(-1)^{k+1}}{R_{f,t \rightarrow N}^k} \left(\tau \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)} - \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)}[\tilde{Q}_{t,\tau}] \right)}{1 + \sum_{k=1}^3 \frac{(-1)^{k+1}}{R_{f,t \rightarrow N}^k} \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)}} \right). \quad \blacksquare \end{aligned}$$

Proposition 6.4 provides a bound on the physical CDF that requires no parameter estimation and relies solely on time t information. This result complements recent work on belief recovery. Ross (2015) demonstrated CDF recovery under the assumption of transition independence, but subsequent research has questioned this

²⁶Chabi-Yo and Loudis (2020) argue that all odd risk-neutral moments should be negative, since they expose the investor to unfavorable market conditions.

assumption (Borovička et al., 2016; Qin et al., 2018; Jackwerth and Menner, 2020). In contrast, Proposition 6.4 establishes a lower bound on the left-tail of the physical distribution using a different set of mild economic constraints. Additionally, Section 2.3 showed that the right-tail of F_t can be approximately recovered from the risk-neutral distribution due to the minimal need for risk-adjustment. These findings suggest the potential for approximate recovery of F_t using option prices.

I will test this hypothesis using the lower bound on disaster risk premia in Proposition 6.5. Specifically, a tight lower bound in (6.9) would enable direct inference on both the physical distribution ($Q_{t,\tau}$) and disaster risk premia ($Q_{t,\tau} - \tilde{Q}_{t,\tau}$). While Section 2.2.1 proposed the quantile model (2.3) to estimate $Q_{t,\tau}$, it can be criticized for having time-homogeneous coefficients. Proposition 6.5 relaxes that assumption. Furthermore, the lower bound in (6.9) is not prone to the historical sample bias critique of Welch and Goyal (2008). Alternatively, one can estimate a disaster risk model to infer $Q_{t,\tau}$, but this approach is also susceptible to misspecification concerns and faces challenges in estimation due to the scarcity of disaster events in the data (Julliard and Ghosh, 2012; Martin, 2013).

6.3 Calculating the Lower Bound

Before assessing how tight the lower bound is in Proposition 6.5, I outline the procedure to calculate it, which depends on $\text{CLB}_{t,\tau}$ and $\tilde{f}_t(\tilde{Q}_{t,\tau})$. Both functions can be derived from $\tilde{Q}_{t,\tau}$, which is estimated using the same data and procedure of Section 2.3. To see that $\tilde{f}_t(\tilde{Q}_{t,\tau})$ can be derived from $\tilde{Q}_{t,\tau}$, notice that $\frac{d}{d\tau}\tilde{Q}_t(\tau) = 1/\tilde{f}_t(\tilde{Q}_{t,\tau})$. The latter term can thus be approximated by²⁷

$$\frac{1}{\tilde{f}_t(\tilde{Q}_{t,\tau})} \approx \frac{\tilde{Q}_t(\tau + h) - \tilde{Q}_t(\tau - h)}{2h},$$

where h is the bandwidth of the τ -grid. Second, to calculate $\text{CLB}_{t,\tau}$ in (6.8), I use $\tilde{Q}_{t,\tau}$, as well as the formula for high-order risk-neutral moments in Appendix A.11.

Given the evidence in Table 2 that $Q_{t,\tau} > \tilde{Q}_{t,\tau}$ in the left-tail, Proposition 6.5 has nontrivial content in the data if $\text{LB}_{t,\tau} \geq 0$. Appendix Table F2 contains summary statistics of $\text{LB}_{t,\tau}$, which show that the lower bound is always positive, right-skewed, more pronounced in the right-tail and economically meaningful in magnitude, with outliers that can spike up to 29%.

²⁷I slightly abuse notation to emphasize that the derivative is taken w.r.t. τ , so that $\tilde{Q}_t(\tau + h)$ denotes $\tilde{Q}_{t,\tau+h}$.

6.4 Tightness of the Lower Bound: In-sample Evidence

To test whether the lower bound in Proposition 6.5 is tight, I form *excess quantile returns*: $R_{m,t \rightarrow N} - \tilde{Q}_{t,\tau}$. Since $\tilde{Q}_{t,\tau}$ is observed at time t , it follows that $Q_{t,\tau}(R_{m,t \rightarrow N} - \tilde{Q}_{t,\tau}) = Q_{t,\tau}(R_{m,t \rightarrow N}) - \tilde{Q}_{t,\tau}$. Subsequently, I use QR to estimate the model

$$Q_{t,\tau}(R_{m,t \rightarrow N}) - \tilde{Q}_{t,\tau}(R_{m,t \rightarrow N}) = \beta_0(\tau) + \beta_1(\tau)\text{LB}_{t,\tau},$$

$$[\hat{\beta}_0(\tau), \hat{\beta}_1(\tau)] = \arg \min_{(\beta_0, \beta_1) \in \mathbb{R}^2} \sum_{t=1}^T \rho_\tau(R_{m,t \rightarrow N} - \tilde{Q}_{t,\tau} - \beta_0 - \beta_1 \text{LB}_{t,\tau}). \quad (6.10)$$

Regression (6.10) is a quantile analogue of the mean excess return regressions of Welch and Goyal (2008). Under the null hypothesis that the lower bound is tight, it holds that

$$H_0 : [\beta_0(\tau), \beta_1(\tau)] = [0, 1]. \quad (6.11)$$

Less restrictive, one can test whether $\beta_0(\tau) = 0$ and $\beta_1(\tau) > 0$, which implies that the statistical “factor” $\text{LB}_{t,\tau}$ explains the conditional quantile wedge.²⁸

Table 5 presents the results of regression (6.10). The null hypothesis of a tight lower bound in (6.11) is not rejected for $\tau = 0.2$, but it is rejected for $\tau \in \{0.05, 0.1\}$ across all horizons. When the null hypothesis is rejected, the $\beta_1(\tau)$ -coefficient exceeds 1, consistent with the theory that $\text{LB}_{t,\tau}$ represents a lower bound on disaster risk premia. In all cases, the lower bound is economically meaningful, since $\beta_1(\tau)$ is significantly different from 0, while $\beta_0(\tau) = 0$ can never be rejected. The explanatory power of the regression in Table 5 is modest, as shown by the $R^1(\tau)$ measure-of-fit:

$$R^1(\tau) = 1 - \frac{\min_{b_0, b_1} \sum \rho_\tau(R_{m,t \rightarrow N} - b_0 - b_1 \text{LB}_{t,\tau})}{\min_{b_0} \sum \rho_\tau(R_{m,t \rightarrow N} - b_0)}. \quad (6.12)$$

But, following the reasoning of Section 5.2, the predictive power in the left-tail cannot be too big, for otherwise near-arbitrage opportunities exist.

I also directly test the predictive power of the lower bound in estimating the physical quantile function. To this end, I use the following model-free quantile forecast:

$$\hat{Q}_{t,\tau} := \tilde{Q}_{t,\tau} + \text{LB}_{t,\tau}. \quad (6.13)$$

²⁸For example, if we start with a quantile factor model $Q_{t,\tau} = \tilde{Q}_{t,\tau} + \beta(\tau)\text{LB}_{t,\tau}$, the model has one testable implication for the data: the intercept in a quantile regression of $R_{m,t \rightarrow N} - \tilde{Q}_{t,\tau}$ on $\text{LB}_{t,\tau}$ should be zero. Quantile factor models have recently been proposed by Chen et al. (2021).

Table 5: **Quantile regression with lower bound**

Horizon	τ	$\hat{\beta}_0(\tau)$	$\hat{\beta}_1(\tau)$	Wald test (p -value)	$R^1(\tau)[\%]$	Obs
<u>30 days</u>	0.05	-0.01 (0.005)	4.43 (0.349)	0.00	6.03	4333
	0.1	-0.01 (0.006)	2.17 (0.450)	0.03	3.18	
	0.2	-0.01 (0.006)	1.33 (0.400)	0.02	0.41	
<u>60 days</u>	0.05	-0.01 (0.013)	5.53 (0.571)	0.00	3.60	4312
	0.1	-0.02 (0.011)	3.25 (0.540)	0.00	2.23	
	0.2	-0.02 (0.009)	1.50 (0.398)	0.27	0.48	
<u>90 days</u>	0.05	-0.02 (0.032)	6.37 (1.113)	0.00	4.91	4291
	0.1	-0.02 (0.018)	3.05 (0.528)	0.00	4.43	
	0.2	-0.02 (0.019)	1.36 (0.626)	0.69	1.46	

Note: This table reports the QR estimates of (6.10) over the sample period 2003-2021 at different horizons, using overlapping returns. Standard errors are shown in parentheses and calculated using SETBB with a block length equal to the prediction horizon. *Wald test* denotes the p -value of the joint restriction $[\beta_0(\tau), \beta_1(\tau)] = [0, 1]$. $R^1(\tau)$ denotes the goodness-of-fit measure (6.12).

To evaluate the accuracy of this forecast, I use QR to estimate the model

$$Q_{t,\tau} = \beta_0(\tau) + \beta_1(\tau)\hat{Q}_{t,\tau}. \quad (6.14)$$

An accurate forecast would imply the joint restriction

$$H_0 : \beta_0(\tau) = 0, \quad \beta_1(\tau) = 1. \quad (6.15)$$

Table 6 summarizes the estimates of (6.14) for several percentiles. The results compare favorably to the risk-neutral estimates in Table 2. First, the point estimates are closer to the $[0, 1]$ benchmark. Second, the Wald test on the joint restriction in (6.15) is never rejected except for $\tau = 0.05$ at the 60-day horizon. Third, the in-sample explanatory power is higher. The same conclusion applies when comparing the predictive results to the expanding quantile regression from Table B1 in the Appendix. Collectively, these findings suggest that $\hat{Q}_{t,\tau}$ can be considered as a good lower bound on the physical quantile function in the left-tail.

Table 6: **Quantile regression with model-free quantile forecast**

Horizon	τ	$\hat{\beta}_0(\tau)$	$\hat{\beta}_1(\tau)$	Wald test (p -value)	$R^1(\tau)[\%]$	$R_{oos}^1(\tau)[\%]$	Obs
30 days	0.05	0.29 (0.249)	0.70 (0.265)	0.06	6.28	9.94	4333
	0.1	0.28 (0.250)	0.72 (0.260)	0.18	3.57	4.02	
	0.2	0.57 (0.381)	0.43 (0.388)	0.29	0.58	2.53	
60 days	0.05	0.30 (0.382)	0.71 (0.426)	0.02	3.40	17.81	4312
	0.1	0.38 (0.352)	0.61 (0.373)	0.13	2.35	9.22	
	0.2	0.44 (0.487)	0.56 (0.498)	0.21	0.57	4.28	
90 days	0.05	0.36 (0.520)	0.64 (0.602)	0.05	4.26	21.98	4291
	0.1	0.31 (0.482)	0.70 (0.521)	0.06	4.19	13.22	
	0.2	0.23 (0.696)	0.78 (0.710)	0.48	0.70	5.99	

Note: This table reports the QR estimates of (6.14) over the sample period 2003-2021. Standard errors are shown in parentheses and calculated using the SETBB, with block length equal to the prediction horizon. *Wald test* gives the p -value of the Wald test on the joint restriction: $\hat{\beta}_0(\tau) = 0, \hat{\beta}_1(\tau) = 1$. $R^1(\tau)$ denotes the in-sample goodness-of fit criterion (2.5). $R_{oos}^1(\tau)$ is the out-of-sample goodness-of fit, using a rolling window size equal to 10 times the return horizon.

6.5 Tightness of the Lower Bound: Out-of-sample Evidence

Given that the in-sample results from Table 6 suggest that $\hat{Q}_{t,\tau}$ is a good lower bound for $Q_{t,\tau}$, it is natural to assess its out-of-sample performance by using $\hat{Q}_{t,\tau}$ to directly predict $Q_{t,\tau}$, which does not require any parameter estimation.

To assess the out-of-sample performance, I use the $R_{oos}^1(\tau)$ measure of fit defined in (2.6) with $\hat{Q}_{t,\tau}$ instead of $\tilde{Q}_{t,\tau}$. Table 6 shows that $\hat{Q}_{t,\tau}$ improves upon the historical rolling quantile out-of-sample in all cases. In particular, this outperformance is most pronounced at the 5th percentile, which is expected since option data are known to provide useful information about extreme downfalls in the stock market (Bates, 2008; Bollerslev and Todorov, 2011). In Appendix F.2, I run a battery of robustness tests which show that, out-of-sample, $LB_{t,\tau}$ better predicts the conditional quantile function than other benchmarks such as the risk-neutral quantile or the VIX index. The latter result is particularly encouraging since the VIX predictor uses in-sample information.

6.6 Robustness of the QR Estimates

The in- and out-of-sample results support $LB_{t,\tau}$ as a robust lower bound for disaster risk premia. It is instructive to compare this lower bound to the disaster risk premia reported in Figure 3b, which are inferred from the quantile regression in (2.4). Consistent with the theory, the estimated disaster risk premium at $\tau = 0.05$ exceeds the lower bound in 99% of cases for 30-day returns and 99.9% for 60-day returns. When violations of the lower bound occur, the differences are typically small.

Figures 8a and 8b show the lower bound for 30- and 60-day returns, respectively, alongside the disaster risk premium estimated from the quantile regression (2.4). In both cases, there is a substantial correlation between the lower bound and the disaster risk premium obtained from the QR estimates. Especially during the global financial crisis and the Covid-19 crisis, both methods predict significant increases in the disaster risk premium. Outside these crisis periods, the lower bound is more conservative. Overall, the model-free lower bound corroborates the robustness of the estimated disaster risk premium in Figure 3b.




Figure 8: **Lower bound on disaster risk premium at 5th percentile.** Panel (a) shows the lower bound on the disaster risk premium for 30-day returns, at $\tau = 0.05$. QR denotes the estimated disaster risk premium from the quantile regression (2.4). The right panel shows a similar graph for 60-day returns. Both figures are smoothed using a 30-day rolling window. The two shaded bars signify the Great Recession period (Dec 2007 – June 2009) and Covid-19 crisis (Feb 2020 – April 2020).

7 Conclusion

I use return and option data on the S&P500 in combination with quantile regression to estimate local differences between the conditional risk-neutral and physical quantile functions. Empirically, these differences are substantial in the left-tail, whereas

in the right-tail, they are barely discernible. Therefore, the lion's share of the equity premium is driven by downside returns, which is model-free evidence for disaster risk.

By tracking these quantile differences over time, the results also demonstrate that disaster risk is time-varying, pervasive, and a driving force behind much of the equity premium, even outside crisis periods. Additionally, my findings show that disaster risk is more nuanced than previous literature suggests. Much of the disagreement can be attributed to the incorporation of conditioning information. While prior research primarily focused on unconditional estimation, my approach accounts for conditioning information embedded in the risk-neutral quantile function, which is crucial to obtain accurate estimates of disaster risk.

To build on this finding, I show that disaster risk makes the SDF highly volatile. In particular, option strategies involving a short position in an asset that pays one dollar in case of a disaster exhibit substantially higher Sharpe ratios compared to a direct investment in the market portfolio. The data reveal that such investment strategies yield a monthly Sharpe ratio of 30%, more than doubling the Sharpe ratio of the market return.

Finally, I suggest a model-free lower bound on disaster risk premia observed from option prices. This lower bound serves as a good predictor of the quantile wedge, exhibiting spikes during crises and significant fluctuations over time. Furthermore, the lower bound closely aligns with estimates of disaster risk premia based on quantile regression, thereby reinforcing the robustness of my findings.

References

- Hengjie Ai and Anmol Bhandari. Asset pricing with endogenously uninsurable tail risk. *Econometrica*, 89(3):1471–1505, 2021. doi:[10.3982/ECTA15142](https://doi.org/10.3982/ECTA15142).
- Yacine Aït-Sahalia and Andrew W. Lo. Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance*, 53(2):499–547, 1998. doi:[10.1111/0022-1082.215228](https://doi.org/10.1111/0022-1082.215228).
- Yacine Aït-Sahalia and Andrew W. Lo. Nonparametric risk management and implied risk aversion. *Journal of Econometrics*, 94(1):9–51, 2000. doi:[10.1016/S0304-4076\(99\)00016-0](https://doi.org/10.1016/S0304-4076(99)00016-0).
- Caio Almeida and René Garcia. Assessing misspecified asset pricing models with empirical likelihood estimators. *Journal of Econometrics*, 170(2):519–537, 2012. doi:[10.1016/j.jeconom.2012.05.020](https://doi.org/10.1016/j.jeconom.2012.05.020).

- Fernando Alvarez and Urban J. Jermann. Using asset prices to measure the persistence of the marginal utility of wealth. *Econometrica*, 73(6):1977–2016, 2005. doi:[10.1111/j.1468-0262.2005.00643.x](https://doi.org/10.1111/j.1468-0262.2005.00643.x).
- Torben G. Andersen, Oleg Bondarenko, Viktor Todorov, and George Tauchen. The fine structure of equity-index option dynamics. *Journal of Econometrics*, 187(2): 532–546, 2015. doi:[10.1016/j.jeconom.2015.02.037](https://doi.org/10.1016/j.jeconom.2015.02.037).
- Joshua Angrist, Victor Chernozhukov, and Iván Fernández-Val. Quantile regression under misspecification, with an application to the U.S. wage structure. *Econometrica*, 74(2):539–563, 2006. doi:[10.1111/j.1468-0262.2006.00671.x](https://doi.org/10.1111/j.1468-0262.2006.00671.x).
- David Backus, Mikhail Chernov, and Ian Martin. Disasters implied by equity index options. *Journal of Finance*, 66(6):1969–2012, 2011. doi:[10.1111/j.1540-6261.2011.01697.x](https://doi.org/10.1111/j.1540-6261.2011.01697.x).
- David Backus, Mikhail Chernov, and Stanley Zin. Sources of entropy in representative agent models. *Journal of Finance*, 69(1):51–99, 2014. doi:[10.1111/jofi.12090](https://doi.org/10.1111/jofi.12090).
- Gurdip Bakshi, Dilip Madan, and George Panayotov. Returns of claims on the upside and the viability of U-shaped pricing kernels. *Journal of Financial Economics*, 97(1):130–154, 2010. doi:[10.1016/j.jfineco.2010.03.009](https://doi.org/10.1016/j.jfineco.2010.03.009).
- Gurdip Bakshi, Fousseni Chabi-Yo, and Xiaohui Gao. A recovery that we can trust? Deducing and testing the restrictions of the recovery theorem. *Review of Financial Studies*, 31(2):532–555, 2018. doi:[10.1093/rfs/hhx108](https://doi.org/10.1093/rfs/hhx108).
- Ravi Bansal and Bruce N. Lehmann. Growth-optimal portfolio restrictions on asset pricing models. *Macroeconomic Dynamics*, 1(2):333–354, 1997. doi:[10.1017/S1365100597003039](https://doi.org/10.1017/S1365100597003039).
- Ravi Bansal and Amir Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance*, 59(4):1481–1509, 2004. doi:[10.1111/j.1540-6261.2004.00670.x](https://doi.org/10.1111/j.1540-6261.2004.00670.x).
- Andrea Barletta and Paolo Santucci de Magistris. Analyzing the risks embedded in option prices with *rndfittool*. *Risks*, 6(2):1–15, 2018. doi:[10.3390/risks6020028](https://doi.org/10.3390/risks6020028).
- Robert J. Barro. Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics*, 121(3):823–866, 2006. doi:[10.1162/qjec.121.3.823](https://doi.org/10.1162/qjec.121.3.823).
- Robert J. Barro. Rare disasters, asset prices, and welfare costs. *American Economic Review*, 99(1):243–264, 2009. doi:[10.1257/aer.99.1.243](https://doi.org/10.1257/aer.99.1.243).

- David S. Bates. The crash of '87: was it expected? The evidence from options markets. *Journal of Finance*, 46(3):1009–1044, 1991. doi:[10.1111/j.1540-6261.1991.tb03775.x](https://doi.org/10.1111/j.1540-6261.1991.tb03775.x).
- David S. Bates. Post-'87 crash fears in the S&P500 futures option market. *Journal of Econometrics*, 94(1):181–238, 2000. doi:[10.1016/S0304-4076\(99\)00021-4](https://doi.org/10.1016/S0304-4076(99)00021-4).
- David S. Bates. The market for crash risk. *Journal of Economic Dynamics and Control*, 32(7):2291–2321, 2008. doi:[10.1016/j.jedc.2007.09.020](https://doi.org/10.1016/j.jedc.2007.09.020).
- Brendan K. Beare and Lawrence D. W. Schmidt. An empirical test of pricing kernel monotonicity. *Journal of Applied Econometrics*, 31(2):338–356, 2016. doi:[10.1002/jae.2422](https://doi.org/10.1002/jae.2422).
- Tyler Beason and David Schreindorfer. Dissecting the equity premium. *Journal of Political Economy*, 130(8):2203–2222, 2022. doi:[10.1086/720396](https://doi.org/10.1086/720396).
- Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973. doi:[10.1086/260062](https://doi.org/10.1086/260062).
- Tim Bollerslev and Viktor Todorov. Tails, fears, and risk premia. *Journal of Finance*, 66(6):2165–2211, 2011. doi:[10.1111/j.1540-6261.2011.01695.x](https://doi.org/10.1111/j.1540-6261.2011.01695.x).
- Tim Bollerslev, Viktor Todorov, and Lai Xu. Tail risk premia and return predictability. *Journal of Financial Economics*, 118(1):113–134, 2015. doi:[10.1016/j.jfineco.2015.02.010](https://doi.org/10.1016/j.jfineco.2015.02.010).
- Jaroslav Borovička, Lars Peter Hansen, and José A. Scheinkman. Misspecified recovery. *Journal of Finance*, 71(6):2493–2544, 2016. doi:[10.1111/jof.12404](https://doi.org/10.1111/jof.12404).
- Douglas T. Breeden and Robert H. Litzenberger. Prices of state-contingent claims implicit in option prices. *Journal of Business*, 51(4):621–651, 1978. doi:[10.1086/296025](https://doi.org/10.1086/296025).
- Mark Broadie, Mikhail Chernov, and Michael Johannes. Understanding index option returns. *Review of Financial Studies*, 22(11):4493–4529, 2009. doi:[10.1093/rfs/hhp032](https://doi.org/10.1093/rfs/hhp032).
- John Y. Campbell and John H. Cochrane. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251, 1999.
- John Y. Campbell and Samuel B. Thompson. Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies*, 21(4):1509–1531, 2008. doi:[10.1093/rfs/hhm055](https://doi.org/10.1093/rfs/hhm055).

- George Casella and Roger L. Berger. *Statistical Inference*. Duxbury/Thomson Learning. Belmont, CA, 2002.
- Fousseni Chabi-Yo and Johnathan Loudis. The conditional expected market return. *Journal of Financial Economics*, 137(3):752–786, 2020. doi:[10.1016/j.jfineco.2020.03.009](https://doi.org/10.1016/j.jfineco.2020.03.009).
- Fousseni Chabi-Yo and Johnathan A. Loudis. A decomposition of conditional risk premia and implications for representative agent models. *Management Science*, 2023. doi:[10.1287/mnsc.2022.01663](https://doi.org/10.1287/mnsc.2022.01663).
- Fousseni Chabi-Yo, René Garcia, and Eric Renault. State dependence can explain the risk aversion puzzle. *Review of Financial Studies*, 21(2):973–1011, 2008. doi:[10.1093/rfs/hhm070](https://doi.org/10.1093/rfs/hhm070).
- Liang Chen, Juan J. Dolado, and Jesús Gonzalo. Quantile factor models. *Econometrica*, 89(2):875–910, 2021. doi:[10.3982/ECTA15746](https://doi.org/10.3982/ECTA15746).
- Victor Chernozhukov, Iván Fernández-Val, and Siyi Luo. Distribution regression with sample selection, with an application to wage decompositions in the UK, 2018. URL <https://arxiv.org/abs/1811.11603>.
- John H. Cochrane. *Asset Pricing*. Princeton University Press, 2005.
- George M. Constantinides and Anisha Ghosh. Asset pricing with countercyclical household consumption risk. *Journal of Finance*, 72(1):415–460, 2017. doi:[10.1111/jofi.12471](https://doi.org/10.1111/jofi.12471).
- Joshua D. Coval and Tyler Shumway. Expected option returns. *Journal of Finance*, 56(3):983–1009, 2001. doi:[10.1111/0022-1082.00352](https://doi.org/10.1111/0022-1082.00352).
- Horatio Cuesdeanu and Jens Carsten Jackwerth. The pricing kernel puzzle in forward looking data. *Review of Derivatives Research*, 21(3):253–276, 2018. doi:[10.1007/s11147-017-9140-8](https://doi.org/10.1007/s11147-017-9140-8).
- Jon Danielsson and Casper G. de Vries. Value-at-risk and extreme returns. *Annales d'Économie et de Statistique*, (60):239–270, 2000. doi:[10.2307/20076262](https://doi.org/10.2307/20076262).
- Francis X. Diebold and Robert S. Mariano. Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3):134–144, 1995. doi:[10.1198/073500102753410444](https://doi.org/10.1198/073500102753410444).
- Itamar Drechsler and Amir Yaron. What’s vol got to do with it. *Review of Financial Studies*, 24(1):1–45, 2011. doi:[10.1093/rfs/hhq085](https://doi.org/10.1093/rfs/hhq085).

- Louis Eeckhoudt and Harris Schlesinger. Putting risk in its proper place. *American Economic Review*, 96(1):280–289, 2006. doi:[10.1257/000282806776157777](https://doi.org/10.1257/000282806776157777).
- Robert F. Engle and Simone Manganelli. CAViaR: Conditional autoregressive value at risk by regression quantiles. *Journal of Business & Economic Statistics*, 22(4): 367–381, 2004. doi:[10.1198/073500104000000370](https://doi.org/10.1198/073500104000000370).
- Emmanuel Farhi and François Gourio. Accounting for macro-finance trends: Market power, intangibles, and risk premia. *Brookings Papers on Economic Activity*, (2): 147–250, 2018. doi:[10.1353/eca.2018.0024](https://doi.org/10.1353/eca.2018.0024).
- Stephen Figlewski. Estimating the implied risk-neutral density for the US market portfolio. In *Volatility and Time Series Econometrics: Essays in Honor of Robert Engle*, chapter 15, pages 323–353. Oxford University Press, 03 2010. doi:[10.1093/acprof:oso/9780199549498.003.0015](https://doi.org/10.1093/acprof:oso/9780199549498.003.0015).
- Damir Filipović, Eberhard Mayerhofer, and Paul Schneider. Density approximations for multivariate affine jump-diffusion processes. *Journal of Econometrics*, 176(2): 93–111, 2013. doi:[10.1016/j.jeconom.2012.12.003](https://doi.org/10.1016/j.jeconom.2012.12.003).
- Xavier Gabaix. Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *Quarterly Journal of Economics*, 127(2):645–700, 2012. doi:[10.1093/qje/qjs001](https://doi.org/10.1093/qje/qjs001).
- Karl B. Gregory, Soumendra N. Lahiri, and Daniel J. Nordman. A smooth block bootstrap for quantile regression with time series. *Annals of Statistics*, 46(3): 1138–1166, 2018. doi:[10.1214/17-AOS1580](https://doi.org/10.1214/17-AOS1580).
- Lars Peter Hansen and Robert J. Hodrick. Forward exchange rates as optimal predictors of future spot rates: An econometric analysis. *Journal of Political Economy*, 88(5):829–853, 1980. doi:[10.1086/260910](https://doi.org/10.1086/260910).
- Lars Peter Hansen and Ravi Jagannathan. Implications of security market data for models of dynamic economies. *Journal of Political Economy*, 99(2):225–262, 1991. doi:[10.1086/261749](https://doi.org/10.1086/261749).
- Fushing Hsieh and Bruce W. Turnbull. Nonparametric and semiparametric estimation of the receiver operating characteristic curve. *Annals of Statistics*, 24(1): 25–40, 1996. doi:[10.1214/aos/1033066197](https://doi.org/10.1214/aos/1033066197).
- Elton P. Hsu and S. R. Srinivasa Varadhan. *Probability Theory and Applications*. American Mathematical Society, 1999.
- Marlène Isoré and Urszula Szczerbowicz. Disaster risk and preference shifts in a New Keynesian model. *Journal of Economic Dynamics and Control*, 79:97–125, 2017. doi:[10.1016/j.jedc.2017.04.001](https://doi.org/10.1016/j.jedc.2017.04.001).

- Jens Carsten Jackwerth. Recovering risk aversion from option prices and realized returns. *Review of Financial Studies*, 13(2):433–451, 2000. doi:[10.1093/rfs/13.2.433](https://doi.org/10.1093/rfs/13.2.433).
- Jens Carsten Jackwerth and Marco Menner. Does the Ross recovery theorem work empirically? *Journal of Financial Economics*, 137(3):723–739, 2020. doi:[10.1016/j.jfineco.2020.03.006](https://doi.org/10.1016/j.jfineco.2020.03.006).
- Christian Julliard and Anisha Ghosh. Can rare events explain the equity premium puzzle? *Review of Financial Studies*, 25(10):3037–3076, 2012. doi:[10.1093/rfs/hhs078](https://doi.org/10.1093/rfs/hhs078).
- Roger Koenker and Gilbert Bassett. Regression quantiles. *Econometrica*, 46(1):33–50, 1978. doi:[10.2307/1913643](https://doi.org/10.2307/1913643).
- Roger Koenker and José A. F. Machado. Goodness of fit and related inference processes for quantile regression. *Journal of the American Statistical Association*, 94(448):1296–1310, 1999. doi:[10.1080/01621459.1999.10473882](https://doi.org/10.1080/01621459.1999.10473882).
- Roger W. Lee. The moment formula for implied volatility at extreme strikes. *Mathematical Finance*, 14(3):469–480, 2004. doi:[10.1111/j.0960-1627.2004.00200.x](https://doi.org/10.1111/j.0960-1627.2004.00200.x).
- Erich L. Lehmann. Some concepts of dependence. *Annals of Mathematical Statistics*, 37(5):1137–1153, 1966. doi:[10.1214/aoms/1177699260](https://doi.org/10.1214/aoms/1177699260).
- Matthew Linn, Sophie Shive, and Tyler Shumway. Pricing kernel monotonicity and conditional information. *Review of Financial Studies*, 31(2):493–531, 2018. doi:[10.1093/rfs/hhx095](https://doi.org/10.1093/rfs/hhx095).
- Yan Liu. Index option returns and generalized entropy bounds. *Journal of Financial Economics*, 139(3):1015–1036, 2021. doi:[10.1016/j.jfineco.2020.08.011](https://doi.org/10.1016/j.jfineco.2020.08.011).
- Ian Martin. Consumption-based asset pricing with higher cumulants. *Review of Economic Studies*, 80(2):745–773, 2013. doi:[10.1093/restud/rds029](https://doi.org/10.1093/restud/rds029).
- Ian Martin. What is the expected return on the market? *Quarterly Journal of Economics*, 132(1):367–433, 2017. doi:[10.1093/qje/qjw034](https://doi.org/10.1093/qje/qjw034).
- Ian Martin and Christian Wagner. What is the expected return on a stock? *Journal of Finance*, 74(4):1887–1929, 2019. doi:[10.1111/jofi.12778](https://doi.org/10.1111/jofi.12778).
- Alexander J. McNeil, Rüdiger Frey, and Paul Embrechts. *Quantitative Risk Management: Concepts, Techniques and Tools*. Princeton University Press, 2015.
- Dimitris N. Politis and Joseph P. Romano. The stationary bootstrap. *Journal of the American Statistical Association*, 89(428):1303–1313, 1994. doi:[10.1080/01621459.1994.10476870](https://doi.org/10.1080/01621459.1994.10476870).

- Likuan Qin and Vadim Linetsky. Long-term risk: A martingale approach. *Econometrica*, 85(1):299–312, 2017. doi:[10.3982/ECTA13438](https://doi.org/10.3982/ECTA13438).
- Likuan Qin, Vadim Linetsky, and Yutian Nie. Long forward probabilities, recovery, and the term structure of bond risk premiums. *Review of Financial Studies*, 31(12):4863–4883, 2018. doi:[10.1093/rfs/hhy042](https://doi.org/10.1093/rfs/hhy042).
- Thomas A. Rietz. The equity risk premium a solution. *Journal of Monetary Economics*, 22(1):117–131, 1988. doi:[10.1016/0304-3932\(88\)90172-9](https://doi.org/10.1016/0304-3932(88)90172-9).
- Joshua V. Rosenberg and Robert F. Engle. Empirical pricing kernels. *Journal of Financial Economics*, 64(3):341–372, 2002. doi:[10.1016/S0304-405X\(02\)00128-9](https://doi.org/10.1016/S0304-405X(02)00128-9).
- Stephen A. Ross. *Neoclassical Finance*. Princeton University Press, 2005.
- Stephen A. Ross. The recovery theorem. *Journal of Finance*, 70(2):615–648, 2015. doi:[10.1111/jofi.12092](https://doi.org/10.1111/jofi.12092).
- Paul Schneider. An anatomy of the market return. *Journal of Financial Economics*, 132(2):325–350, 2019. doi:[10.1016/j.jfineco.2018.10.015](https://doi.org/10.1016/j.jfineco.2018.10.015).
- Paul Schneider and Fabio Trojani. (Almost) model-free recovery. *Journal of Finance*, 74(1):323–370, 2019. doi:[10.1111/jofi.12737](https://doi.org/10.1111/jofi.12737).
- David Schreindorfer. Macroeconomic tail risks and asset prices. *Review of Financial Studies*, 33(8):3541–3582, 2020. doi:[10.1093/rfs/hhz105](https://doi.org/10.1093/rfs/hhz105).
- Sang Byung Seo and Jessica A. Wachter. Option prices in a model with stochastic disaster risk. *Management Science*, 65(8):3449–3469, 2019. doi:[10.1287/mnsc.2017.2978](https://doi.org/10.1287/mnsc.2017.2978).
- Robert J. Serfling. *Approximation Theorems of Mathematical Statistics*. John Wiley & Sons, 2009.
- Karl N. Snow. Diagnosing asset pricing models using the distribution of asset returns. *Journal of Finance*, 46(3):955–983, 1991. doi:[10.1111/j.1540-6261.1991.tb03773.x](https://doi.org/10.1111/j.1540-6261.1991.tb03773.x).
- Michael Stutzer. A Bayesian approach to diagnosis of asset pricing models. *Journal of Econometrics*, 68(2):367–397, 1995. doi:[10.1016/0304-4076\(94\)01656-K](https://doi.org/10.1016/0304-4076(94)01656-K).
- Engin A. Sungur. Dependence information in parameterized copulas. *Communications in Statistics - Simulation and Computation*, 19(4):1339–1360, 1990. doi:[10.1080/03610919008812920](https://doi.org/10.1080/03610919008812920).
- Aad W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 2000.

- Jessica A. Wachter. Can time-varying risk of rare disasters explain aggregate stock market volatility? *Journal of Finance*, 68(3):987–1035, 2013.
doi:[10.1111/jof.12018](https://doi.org/10.1111/jof.12018).
- Ivo Welch and Amit Goyal. A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21(4):1455–1508, 2008.
doi:[10.1093/rfs/hhm014](https://doi.org/10.1093/rfs/hhm014).

A Proofs

This section contains proofs and detailed calculations of results used in the main paper.

A.1 Decomposing the Equity Premium

For any atomless integrable random variable X with CDF $F(\cdot)$ and quantile function $Q = F^{-1}$, we have

$$\mathbb{E}(X) = \int_{\mathbb{R}} x \, dF(x) = \int_0^1 Q(\tau) \, d\tau.$$

The second identity holds by the change of variables formula for the Lebesgue-Stieltjes integral. In case F has a density, the formula follows from a simple substitution $x \rightarrow Q(\tau)$. Hence,

$$\mathbb{E}_t[R_{m,t \rightarrow N}] - R_{f,t \rightarrow N} = \mathbb{E}_t[R_{m,t \rightarrow N}] - \tilde{\mathbb{E}}_t(R_{m,t \rightarrow N}) = \int_0^1 (Q_{t,\tau} - \tilde{Q}_{t,\tau}) \, d\tau.$$

A.2 Stochastic Dominance and Pricing Kernel Monotonicity

In this section I provide more details on the relation between stochastic dominance and pricing kernel monotonicity. To begin with, recall that the physical distribution is first-order stochastically dominant (FOSD) over the risk-neutral distribution if and only if $F_t(x) \leq \tilde{F}_t(x)$, or $\tilde{Q}_{t,\tau} \leq Q_{t,\tau}$. The definition is also equivalent to $F_t(\tilde{Q}_{t,\tau}) \leq \tau$ for all $\tau \in (0, 1)$, which follows from the substitution $x \rightarrow \tilde{Q}_{t,\tau}$.

To see the connection with pricing kernel monotonicity, recall from [Beare and Schmidt \(2016\)](#) that pricing kernel monotonicity is equivalent to $\phi_t(\tau) := F_t(\tilde{Q}_{t,\tau})$ being a convex function for all τ .²⁹ Figure A1 shows two different ordinal dominance curves (ODCs); the blue line corresponds to a situation where FOSD holds and the pricing kernel is monotonic (hence convex), whereas the yellow line shows a scenario where FOSD does not hold and convexity automatically fails. The geometric argument for why non-monotonicity is implied by a failure of FOSD is conveyed by the figure: if FOSD fails, the yellow line must cross the 45-degree line for some $\tau \in (0, 1)$, which automatically implies that the ODC is non-convex since the ODC has to satisfy $\phi_t(1) = 1$, because the physical and risk-neutral measures are equivalent. The proposition below thus follows.

Proposition A.1. *If the pricing kernel is a monotonically decreasing function of the market return, the physical measure first-order stochastically dominates the risk-*

²⁹[Beare and Schmidt \(2016\)](#) actually consider the reverse function $\phi_t(\tau) = \tilde{F}_t(Q_{t,\tau})$, so that pricing kernel monotonicity is equivalent to $\phi_t(\cdot)$ being concave.

neutral measure. Conversely, a violation of FOSD implies a violation of pricing kernel monotonicity.

A violation of FOSD is puzzling from the viewpoint of expected utility maximization. In this framework, the SDF is given by $u'(R_{m,t \rightarrow N})/\mathbb{E}_t(u'(R_{m,t \rightarrow N}))$, where $u(\cdot)$ is a utility function and the initial endowment is normalized to one for simplicity. The following proposition shows that a sufficient (but not necessary) condition for FOSD to hold is that $u'(\cdot)$ is non-increasing; a rather ubiquitous assumption in asset pricing models.

Proposition A.2. *In the expected utility framework, a sufficient condition for the physical measure to first-order stochastically dominate the risk-neutral measure is that $u'(\cdot)$ is non-increasing.*

Proof. Using the SDF to change from physical to risk-neutral measure, it follows that FOSD is equivalent to

$$\begin{aligned} F_t(x) &\leq \tilde{F}_t(x) \\ \iff \mathbb{E}_t[\mathbb{1}(R_{m,t \rightarrow N} \leq x)] &\leq \mathbb{E}_t\left[\frac{u'(R_{m,t \rightarrow N})}{\mathbb{E}_t[u'(R_{m,t \rightarrow N})]} \mathbb{1}(R_{m,t \rightarrow N} \leq x)\right] \\ \iff 0 &\leq \text{COV}_t(\mathbb{1}(R_{m,t \rightarrow N} \leq x), u'(R_{m,t \rightarrow N})). \end{aligned}$$

By Lemma E.1, the covariance above is nonnegative if $u'(\cdot)$ is non-increasing. ■

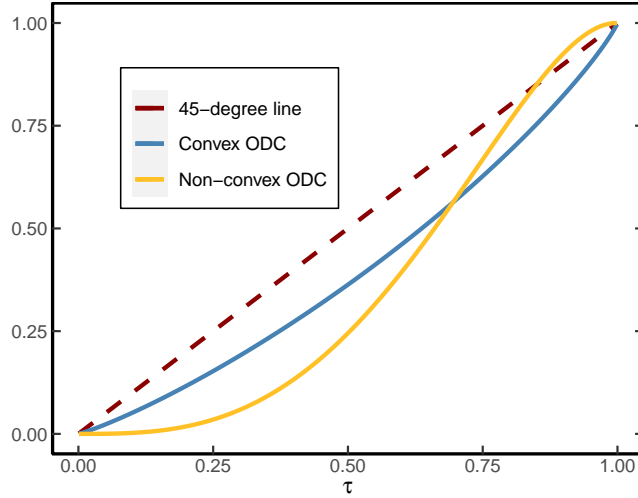


Figure A1: Ordinal dominance curve with and without first-order stochastic dominance. This figure shows two different ordinal dominance curves. The blue ODC corresponds to a situation where the physical measure FOSD the risk-neutral measure, whereas the yellow line shows a situation where FOSD fails.

A.3 Proof of Proposition 4.1

I separately show (i) and (ii) of Proposition 4.1. To prove these results, I use the following lemma.

Lemma A.3. *In the lognormal model, the physical and risk-neutral quantile functions conditional on μ_t, σ_t are given by, respectively*

$$Q_{t,\tau} = \exp \left[\left(\mu_t - \frac{1}{2} \sigma_t^2 \right) N + \sigma_t \sqrt{N} \Phi^{-1}(\tau) \right] \quad (\text{A.1})$$

$$\tilde{Q}_{t,\tau} = \exp \left[\left(r_f - \frac{1}{2} \sigma_t^2 \right) N + \sigma_t \sqrt{N} \Phi^{-1}(\tau) \right], \quad (\text{A.2})$$

where $\Phi^{-1}(\cdot)$ denotes the quantile function of the standard normal distribution. If $\mu_t \sim \mathcal{N}(\mu, \sigma_\mu^2)$ and independent from σ_t , the physical quantile function conditional on σ_t , but not μ_t , equals

$$Q_{t,\tau}(\sigma_t, \sigma_\mu) = \exp \left[\left(\mu - \frac{1}{2} \sigma_t^2 \right) N + \left(\sqrt{\sigma_\mu^2 N^2 + \sigma_t^2 N} \right) \Phi^{-1}(\tau) \right]. \quad (\text{A.3})$$

Proof. The quantile function of a random variable X such that $\log X \sim \mathcal{N}(a, b^2)$, is given by $\exp(a + b\Phi^{-1}(\tau))$. Therefore, the quantile functions conditional on μ_t, σ_t in (A.1) and (A.2) follow immediately from the conditional lognormal assumption. In (A.3), the function is conditioned on σ_t , but not μ_t . Since μ_t is assumed to be normally distributed and independent from σ_t , it follows that

$$\left(\mu_t - \frac{1}{2} \sigma_t^2 \right) N + \sigma_t \sqrt{N} Z_{t+N} | \sigma_t \sim \mathcal{N} \left(\left(\mu - \frac{1}{2} \sigma_t^2 \right) N, \sigma_\mu^2 N^2 + \sigma_t^2 N \right).$$

The expression in (A.3) can now be obtained again using the general formula of the lognormal quantile function. ■

Proof of Proposition 4.1(i). Recall that $\sqrt{a^2 + b^2} \leq \sqrt{a^2 + b^2 + 2ab} = a + b$, provided $a, b \geq 0$. This inequality shows that

$$\begin{aligned} & \exp \left[\left(\sqrt{\sigma_\mu^2 N^2 + \sigma_t^2 N} - \sigma_t \sqrt{N} \right) \Phi^{-1}(\tau) \right] \\ & \leq \exp \left[\left(\sqrt{\sigma_\mu^2 N^2 + \sigma_t^2 N} - \sigma_t \sqrt{N} \right) |\Phi^{-1}(\tau)| \right] \\ & \leq \exp \left(\sigma_\mu N |\Phi^{-1}(\tau)| \right) \\ & = 1 + \mathcal{O}(\sigma_\mu N), \end{aligned}$$

uniformly in $\tau \in \mathcal{I}$ and the support of σ_t . In combination with Lemma A.3, it follows

that

$$\begin{aligned} Q_{t,\tau}(\sigma_t, \sigma_\mu) &= \tilde{Q}_{t,\tau} e^{(\mu-r_f)N} \exp \left[\left(\sqrt{\sigma_\mu^2 N^2 + \sigma_t^2 N} - \sigma_t \sqrt{N} \right) \Phi^{-1}(\tau) \right] \\ &= \tilde{Q}_{t,\tau} e^{(\mu-r_f)N} (1 + \mathcal{O}(\sigma_\mu N)). \end{aligned} \quad \blacksquare$$

In order to prove Proposition 4.1(ii), I need additional regularity conditions stated in Assumption A.4 below. The following notation for the quantile empirical process will be used:

$$\begin{aligned} L_{T,\tau}(\beta, \sigma_\mu) &:= \frac{1}{T} \sum_{t=1}^T \rho_\tau(R_{m,t \rightarrow N} - \beta_0 - \beta_1 \tilde{Q}_{t,\tau}) \\ L_\tau(\beta, \sigma_\mu) &:= \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[\rho_\tau(R_{m,t \rightarrow N} - \beta_0 - \beta_1 \tilde{Q}_{t,\tau}) \right]. \end{aligned}$$

Assumption A.4. *In the lognormal model, assume additionally that*

- (i) $\mathbb{E}[R_{m,t \rightarrow N}]$ and $\mathbb{E}[\tilde{Q}_{t,\tau}]$ are finite,
- (ii) $L_\tau(\beta, 0)$ has an identifiably unique minimum β^* at $\sigma_\mu = 0$, i.e., for all $\varepsilon > 0$

$$\inf_{\|\beta - \beta^*\| > \varepsilon} L_\tau(\beta, 0) - L_\tau(\beta^*, 0) > 0.$$

- (iii) as $T \rightarrow \infty$, for any compact set \mathcal{B} and sequence $b_T \searrow 0$,

$$\sup_{\beta \in \mathcal{B}} \|L_\tau(\beta, \sigma_\mu^T) - L_\tau(\beta, 0)\| = o(1) \text{ (Uniform continuity)}. \quad (\text{A.4a})$$

$$\sup_{\sigma_\mu \leq b_T} \sup_{\beta \in \mathcal{B}} \|L_{T,\tau}(\beta, \sigma_\mu) - L_\tau(\beta, \sigma_\mu)\| = o_p(1) \text{ (Uniform LLN)}. \quad (\text{A.4b})$$

Proof of Proposition 4.1(ii). Consider the population minimization problem of quantile regression at $\sigma_\mu = 0$

$$[\beta_0^*(0; \tau), \beta_1^*(0; \tau)] := \arg \min_{(\beta_0, \beta_1) \in \mathbb{R}^2} L_\tau(\beta, 0). \quad (\text{A.5})$$

Assumptions A.4(i,ii) ensure that the objective function is well defined and the solution in (A.5) is unique for all $\tau \in \mathcal{I}$. At $\sigma_\mu = 0$, $Q_{t,\tau} = e^{(\mu-r)N}$, so that $[\beta_0^*(0; \tau), \beta_1^*(0; \tau)] = [0, e^{(\mu-r)N}]$. To ease notation in the following derivation, I write $\hat{\beta}(\sigma_\mu^T) := \arg \min_\beta L_{T,\tau}(\beta, \sigma_\mu^T)$ and $\beta^*(0) = \arg \min_\beta L_\tau(\beta, 0)$. It then follows

that for every $\varepsilon > 0$ there exists a $\delta > 0$ such that

$$\begin{aligned}
& \mathbb{P} \left(\left\| \widehat{\beta}(\sigma_\mu^T) - \beta^*(0) \right\| > \varepsilon \right) \\
& \leq \mathbb{P} \left(L_\tau(\widehat{\beta}(\sigma_\mu^T), 0) - L_\tau(\beta^*(0), 0) > \delta \right) \\
& = \mathbb{P} \left(L_\tau(\widehat{\beta}(\sigma_\mu^T), 0) - L_{T,\tau}(\widehat{\beta}(\sigma_\mu^T), \sigma_\mu^T) + L_{T,\tau}(\widehat{\beta}(\sigma_\mu^T), \sigma_\mu^T) - L_\tau(\beta^*(0), 0) > \delta \right) \\
& \leq \mathbb{P} \left(L_\tau(\widehat{\beta}(\sigma_\mu^T), 0) - L_{T,\tau}(\widehat{\beta}(\sigma_\mu^T), \sigma_\mu^T) + L_{T,\tau}(\beta^*(0), \sigma_\mu^T) - L_\tau(\beta^*(0), 0) > \delta \right) \\
& \leq \mathbb{P} \left(2 \sup_{\beta \in \mathcal{B}} \|L_\tau(\beta, 0) - L_{T,\tau}(\beta, \sigma_\mu^T)\| > \delta \right).
\end{aligned}$$

The second line follows from identification and the second to last line from the minimization property of $\widehat{\beta}(\sigma_\mu^T)$. Therefore, it suffices to show that

$$\sup_{\beta \in \mathcal{B}} \|L_\tau(\beta, 0) - L_{T,\tau}(\beta, \sigma_\mu^T)\| = o_p(1).$$

This claim follows from

$$\begin{aligned}
& \sup_{\beta \in \mathcal{B}} \|L_\tau(\beta, 0) - L_{T,\tau}(\beta, \sigma_\mu^T)\| \\
& \leq \sup_{\beta \in \mathcal{B}} \|L_\tau(\beta, 0) - L_\tau(\beta, \sigma_\mu^T)\| + \|L_\tau(\beta, \sigma_\mu^T) - L_{T,\tau}(\beta, \sigma_\mu^T)\| \\
& \leq \sup_{\beta \in \mathcal{B}} \|L_\tau(\beta, 0) - L_\tau(\beta, \sigma_\mu^T)\| + \sup_{\sigma_\mu \leq b_T} \sup_{\beta \in \mathcal{B}} \|L_\tau(\beta, \sigma_\mu) - L_{T,\tau}(\beta, \sigma_\mu)\|.
\end{aligned}$$

The first term is $o(1)$ by (A.4a) and the second term is $o_p(1)$ by (A.4b), which completes the proof. The claim in (4.5) easily follows from (4.4). ■

A.4 Proof of Proposition 5.1

Proof. Starting from the definition of the risk-neutral quantile function, it follows that

$$\begin{aligned}
\tau &= \tilde{\mathbb{P}}_t \left[R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right] \\
&= \tilde{\mathbb{E}}_t \left[\mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right) \right] \\
&= \frac{1}{\mathbb{E}_t [M_{t \rightarrow N}]} \mathbb{E}_t \left[M_{t \rightarrow N} \mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right) \right] \\
&= \frac{1}{\mathbb{E}_t [M_{t \rightarrow N}]} \left(\text{COV}_t \left(M_{t \rightarrow N}, \mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right) \right) + \mathbb{E}_t [M_{t \rightarrow N}] \mathbb{E}_t \left[\mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right) \right] \right) \\
&= \frac{1}{\mathbb{E}_t [M_{t \rightarrow N}]} \text{COV}_t \left(M_{t \rightarrow N}, \mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right) \right) + \underbrace{\mathbb{E}_t \left[\mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right) \right]}_{=\phi_t(\tau)}.
\end{aligned} \tag{A.6}$$

Rearranging then yields

$$\frac{1}{\mathbb{E}_t [M_{t \rightarrow N}]} \text{COV}_t \left(M_{t \rightarrow N}, \mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right) \right) = \tau - \phi_t(\tau).$$

Using Cauchy-Schwarz renders the inequality

$$\begin{aligned}
\frac{1}{\mathbb{E}_t [M_{t \rightarrow N}]} \sigma_t(M_{t \rightarrow N}) \sigma_t \left(\mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right) \right) &\geq |\tau - \phi_t(\tau)| \\
\frac{\sigma_t(M_{t \rightarrow N})}{\mathbb{E}_t [M_{t \rightarrow N}]} &\geq \frac{|\tau - \phi_t(\tau)|}{\sigma_t \left(\mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right) \right)}. \tag{A.7}
\end{aligned}$$

Finally, since $\mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right)$ is a Bernoulli random variable, it follows that

$$\sigma_t \left(\mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right) \right) = \sqrt{\phi_t(\tau)(1 - \phi_t(\tau))}. \tag{A.8}$$

Proposition 5.1 now follows after substituting (A.8) into (A.7). The bound formulated in terms of the CDFs in (5.2) follows from the substitution $\tilde{Q}_{t,\tau} \rightarrow x$. ■

A.5 Distribution Bound when SDF and Return are Jointly Normal

In this Section I derive (5.7) and (5.8), when M and R_m are jointly normal. First consider (5.8). The proof of the distribution bound in Proposition 5.1 gives the following identity

$$\frac{|\tau - \phi(\tau)|}{R_f} = \left| \text{COV} \left(\mathbb{1} \left(R_m \leq \tilde{Q}_\tau \right), M \right) \right|.$$

Standard SDF properties also yield the well known result

$$\frac{|\mathbb{E}(R_m) - R_f|}{R_f} = |\text{COV}(R_m, M)|.$$

These results, combined with (5.7) prove (5.8), since

$$\begin{aligned} \frac{\text{HJ bound}}{\text{distribution bound}} &= \frac{\frac{|\mathbb{E}[R_m] - R_f|}{\sigma_R R_f}}{\frac{|\tau - \phi(\tau)|}{\sqrt{\phi(\tau)(1 - \phi(\tau))} R_f}} \\ &\stackrel{(5.7)}{=} \frac{\sqrt{\phi(\tau)(1 - \phi(\tau))}}{\sigma_R f_R(\tilde{Q}_\tau)}, \end{aligned}$$

where $f_R(\tilde{Q}_\tau)$ is the marginal density of R_m .

Finally, I make use of the following covariance identities to prove (5.7).

Lemma A.5 (Hoeffding). *For any square integrable random variable X and Z with marginal CDFs F_X, F_Z and joint CDF $F_{X,Z}$, it holds that*

$$\text{COV}[\mathbb{1}(Z \leq z), X] = - \int_{-\infty}^{\infty} [F_{X,Z}(x, z) - F_X(x)F_Z(z)] dx \quad (\text{A.9})$$

$$\text{COV}[Z, X] = - \int_{-\infty}^{\infty} \text{COV}[\mathbb{1}(Z \leq z), X] dz. \quad (\text{A.10})$$

Proof. See [Lehmann \(1966\)](#). ■

I also need a relation for the bivariate normal distribution. Suppose that X, Z are jointly normal with correlation ρ , mean μ_X, μ_Z and variance σ_X^2, σ_Z^2 , then

$$\frac{\partial \Phi_2(x, z; \rho, \mu_X, \mu_Z, \sigma_X^2, \sigma_Z^2)}{\partial \rho} = \sigma_X \sigma_Z \phi_2(x, z; \rho, \mu_X, \mu_Z, \sigma_X^2, \sigma_Z^2), \quad (\text{A.11})$$

where $\Phi_2(\cdot)$ denotes the bivariate normal CDF and $\phi_2(\cdot)$ denotes the bivariate normal PDF ([Sungur, 1990](#)). We can now prove a covariance identity for jointly normal random variables.

Proposition A.6. *Suppose R_m and M are jointly normal with correlation ρ , then*

$$-\text{COV}[\mathbb{1}(R_m \leq x), M] = \text{COV}[R_m, M] \phi_R(x), \quad (\text{A.12})$$

where $\phi_R(\cdot)$ is the marginal density of R_m .

Proof. To lighten notation, I suppress the dependence on $\mu_R, \mu_M, \sigma_R^2, \sigma_M^2$ in the joint

CDF and PDF. We then have

$$\begin{aligned}
-\mathbb{C}\mathbb{O}\mathbb{V}[\mathbb{1}(R_m \leq x), M] &= \int_{-\infty}^{\infty} \Phi_2(x, m; \rho) - \Phi_2(x, m; 0) dm \\
&= \int_{-\infty}^{\infty} \int_0^{\rho} \sigma_R \sigma_M \phi_2(x, m; y) dy dm \\
&= \sigma_R \sigma_M \rho \phi_R(x) \\
&= \mathbb{C}\mathbb{O}\mathbb{V}[R_m, M] \phi_R(x),
\end{aligned}$$

where, in the first line, I use (A.9) together with $F_R(r)F_M(m) = \Phi_2(r, m; 0)$, the second line follows from (A.11) and the third line follows from Fubini's theorem to swap the order of integration and $\int_{-\infty}^{\infty} \phi_2(x, m; y) dm = \phi_R(x)$. ■

Remark 5. The second covariance identity in (A.10) shows that $\mathbb{C}\mathbb{O}\mathbb{V}[\mathbb{1}(R_m \leq x), M]$ is a measure of local dependence. In case of joint normality (A.12), the weight is given by the marginal PDF. For other distributions, the weighting factor is more complicated, but sometimes can be given an explicit form using a local Gaussian representation (see Chernozhukov et al. (2018)).

A.6 Minimizer of Distribution Bound with Normal SDF

This section shows that the relative efficiency between the HJ bound and distribution bound is minimized when $\tilde{Q}_\tau = \mu_R$. To see this, write $x = \tilde{Q}_\tau$, and use $F(\cdot)$ to denote the physical CDF of R_m . I also drop the R subscript for f to avoid notational clutter. Consider

$$\Gamma(x) = \frac{F(x)(1 - F(x))}{f(x)^2}.$$

Minimizing $\Gamma(x)$ is equivalent to minimizing (5.8) and first order conditions imply that the optimal x^* satisfies

$$[f(x^*) - 2F(x^*)f(x^*)]f(x^*)^2 - 2f(x^*)f'(x^*)[F(x^*)(1 - F(x^*))] = 0. \quad (\text{A.13})$$

Since f, F are the respective PDF and CDF of the normal random variable R_m , it follows that $f'(\mu_R) = 0$ and $F(\mu_R) = 1/2$. As a result, (A.13) holds when $\tilde{Q}_{\tau^*} = x^* = \mu_R$.

A.7 Distribution Bound when SDF and Return are Log-normal

This section provides a closed form approximation for the relative efficiency between the HJ bound and distribution bound under joint lognormality. The result depends on Stein's Lemma (Casella and Berger, 2002, Lemma 3.6.5).³⁰

³⁰I use the form of Stein's Lemma reported in Cochrane (2005, p. 163), which follows from Stein's lemma as reported in Casella and Berger (2002).

Lemma A.7 (Stein's Lemma). *If X_1, X_2 are bivariate normal, $g : \mathbb{R} \rightarrow \mathbb{R}$ is differentiable and $\mathbb{E}[g'(X_1)] < \infty$, then*

$$\mathbb{COV}(g(X_1), X_2) = \mathbb{E}[g'(X_1)] \mathbb{COV}(X_1, X_2).$$

To prove the approximation, we approximate M by a first order Taylor expansion, which gives

$$\widehat{M} = e^{-(r_f + \frac{\sigma_M^2}{2})N} + Z_M \sigma_M \sqrt{N} e^{-(r_f + \frac{\sigma_M^2}{2})N}.$$

Notice that $\widehat{M} = M + o_p(\sqrt{N})$. Consequently, by Stein's Lemma

$$\begin{aligned} \mathbb{COV}(R_m, M) &\approx \mathbb{COV}(R_m, \widehat{M}) = \sigma_M \sqrt{N} e^{-(r_f + \frac{\sigma_M^2}{2})N} \mathbb{COV}(R_m, Z_M) \\ &= \sigma_M \sqrt{N} e^{-(r_f + \frac{\sigma_M^2}{2})N} \mathbb{E} \left[\sigma_R \sqrt{N} \exp \left(\left[\mu_R - \frac{\sigma_R^2}{2} \right] N + \sigma_R \sqrt{N} Z_R \right) \right] \mathbb{COV}(Z_R, Z_M) \\ &= \sigma_M \sigma_R N e^{-(r_f + \frac{\sigma_M^2}{2})N} e^{\mu_R N} \mathbb{COV}(Z_R, Z_M). \end{aligned}$$

By Proposition A.6,

$$\begin{aligned} \mathbb{COV}(\mathbb{1}(\log R_m \leq x), M) &\approx \mathbb{COV}(\mathbb{1}(\log R_m \leq x), \widehat{M}) \\ &= \sigma_M \sqrt{N} e^{-(r_f + \frac{\sigma_M^2}{2})N} \mathbb{COV}(\mathbb{1}(\log R_m \leq x), Z_M) \\ &= \sigma_M \sqrt{N} e^{-(r_f + \frac{\sigma_M^2}{2})N} \mathbb{COV}(\mathbb{1}((\mu_R - \sigma_R^2/2)N + \sigma_R \sqrt{N} Z_R \leq x), Z_M) \\ &= -\sigma_M \sqrt{N} e^{-(r_f + \frac{\sigma_M^2}{2})N} f(x) \mathbb{COV}(Z_R, Z_M). \end{aligned}$$

Here, f is the density of a normal random variable with mean $(\mu_R - \sigma_R^2/2)N$ and variance $N\sigma_R^2$. As a result,

$$\left| \frac{\mathbb{E}[R_m] - e^{Nr_f}}{\tau - \phi(\tau)} \right| \approx \frac{\sigma_R \sqrt{N} e^{\mu_R N}}{f(x)}. \quad (\text{A.14})$$

The same reasoning in Example 5.2 implies that the relative efficiency between the HJ and distribution bound can be approximated by

$$\frac{\text{HJ bound}}{\text{distribution bound}} = \frac{\frac{|\mathbb{E}[R_m] - R_f|}{\sigma(R_m) R_f}}{\frac{|\tau - \phi(\tau)|}{\sqrt{\phi(\tau)(1 - \phi(\tau))} R_f}} \quad (\text{A.15})$$

$$\stackrel{(\text{A.14})}{\approx} \frac{\sqrt{\mathbb{P}(r \leq x) \cdot (1 - \mathbb{P}(r \leq x))}}{\sigma(R_m)} \times \frac{\sigma_R \sqrt{N} e^{\mu_R N}}{f(x)}, \quad (\text{A.16})$$

where $r = \log R$ and $x = \log \tilde{Q}_\tau$. Using the same reasoning as in Example 5.2, the expression on the right hand side of (A.15) is minimized by choosing $x = \log \tilde{Q}_\tau^*$ s.t.

$\mathbb{P}(R_m \leq \tilde{Q}_\tau^*) = 1/2$. In that case the relative efficiency equals

$$\frac{\sqrt{2\pi\sigma_R^2}\sqrt{N}e^{\mu_R N}}{2\sqrt{[\exp(\sigma_R^2 N) - 1]\exp(2\mu_R N)}} = \frac{1}{2}\sqrt{\frac{2\pi\sigma_R^2 N}{\exp(\sigma_R^2 N) - 1}}.$$

A.8 Distribution Bound with Pareto Distribution

This section derives an explicit expression of the distribution bound when the return and SDF follow the Pareto distribution.

Example A.1 (Pareto distribution). Let $U \sim \mathbf{Unif}[0, 1]$ (Uniform distribution on $[0, 1]$) and consider the following specification:

$$M = AU^\alpha, \quad R_m = BU^{-\beta} \quad \text{with} \quad \alpha, \beta, A, B > 0. \quad (\text{A.17})$$

A random variable $X \sim \mathbf{Par}(C, \zeta)$ follows a Pareto distribution with scale parameter $C > 0$ and shape parameter $\zeta > 0$ if the CDF is given by

$$\mathbb{P}(X \leq x) = \begin{cases} 1 - (x/C)^{-\zeta} & x \geq C \\ 0 & x < C. \end{cases}$$

The assumption (A.17) implies that returns follow a Pareto distribution, both under the physical and risk-neutral measures. This fact allows me to obtain an explicit expression for the distribution bound. I summarize these properties in the Proposition below.

Proposition A.8. *Let the SDF and return be given by (A.17). Then,*

- (i) *Under \mathbb{P} , the distribution of returns is Pareto: $R_m \sim \mathbf{Par}\left(B, \frac{1}{\beta}\right)$.*
- (ii) *Under $\tilde{\mathbb{P}}$, the distribution of returns is Pareto: $R_m \sim \mathbf{Par}\left(B, \frac{\alpha+1}{\beta}\right)$.*
- (iii) *The Sharpe ratio on the asset return is given by*

$$\frac{\mathbb{E}[R_m] - R_f}{\sigma(R_m)} = \frac{\frac{B}{1-\beta} - \frac{\alpha+1}{A}}{\sqrt{\frac{B^2}{1-2\beta} - \left(\frac{B}{1-\beta}\right)^2}}. \quad (\text{A.18})$$

- (iv) *The distribution bound is given by*

$$\frac{1}{R_f} \frac{|\tau - \phi(\tau)|}{\sqrt{\phi(\tau)(1 - \phi(\tau))}} = \frac{A}{1 + \alpha} \frac{\left| \tau - 1 + (1 - \tau)^{\frac{1}{\alpha+1}} \right|}{\sqrt{(1 - (1 - \tau)^{\frac{1}{\alpha+1}})(1 - \tau)^{\frac{1}{\alpha+1}}}.$$

- (v) *If $\beta \nearrow \frac{1}{2}$, the HJ bound converges to 0.*

Proof. See the end of this section. ■

Proposition A.8(iv) shows that the distribution bound is independent of the Pareto tail index β . Properties (iv) and (v) provide some intuition when the distribution bound is stronger than the HJ bound. Namely, heavier tails of the distribution of R_m (as measured by β) lead to a lower Sharpe ratio. However, the distribution bound is unaffected by β since it only depends on the tail index α . Therefore, when β gets close to $1/2$, the HJ bound is rather uninformative whereas the distribution bound may fare better. Moreover, no additional restrictions on the parameter space are necessary to calculate the distribution bound, while the HJ bound requires $\beta < 1/2$.³¹

Figure A2 shows two instances of the distribution and HJ bound using different parameter calibrations. Both calibrations are targeted to match an equity premium of 8% and risk-free rate of 0%, but in Panel A2b, the distribution of returns has a fatter tail compared to Panel A2a. In both calibrations, the distribution bound has a range of values for which it is stronger than the HJ bound. In line with Proposition A.8, we see that the range is larger in Panel (b), since the HJ bound is less informative owing to the heavier tails of R_m . However, the distribution bound attains its maximum in the right-tail since that is the region where the physical and risk-neutral measure differ most. This result is inconsistent with the empirical results from Table 2, which indicate that the physical and risk-neutral measure are nearly identical in the right-tail.

Proof of Proposition A.8. (i) The distribution of returns is Pareto, since

$$\begin{aligned}\mathbb{P}(R_m \leq x) &= \mathbb{P}(U^{-\beta} \leq x/B) \\ &= \mathbb{P}\left(U \geq (x/B)^{-\frac{1}{\beta}}\right) = 1 - \left(\frac{x}{B}\right)^{-\frac{1}{\beta}}, \quad x \geq B.\end{aligned}$$

(ii) Since $R_f M$ is the Radon-Nikodym derivative that induces a change of measure

³¹The latter restriction is not unreasonable for asset returns, since typical tail index estimates suggest $\beta \in [1/4, 1/3]$ (Danielsson and de Vries, 2000).




Figure A2: **HJ and distribution bound for heavy tailed returns.** Both panels plot the distribution bound, HJ bound and true SDF volatility for the Pareto model (A.17). In Panel (b), the distribution of returns has a fatter tail compared to Panel (a). Panel (a) uses the parameters $[A, \alpha, B, \beta] = [1.19, 0.19, 0.72, 0.33]$. Panel (b) uses the parameters $[A, \alpha, B, \beta] = [1.11, 0.11, 0.59, 0.45]$. Both calibrations imply an equity premium of 8% and (net) risk-free rate of 0%.

from \mathbb{P} to $\tilde{\mathbb{P}}$, it follows that

$$\begin{aligned}
 \tilde{\mathbb{P}}(R_m \leq x) &= R_f \mathbb{E}[M \mathbb{1}(R_m \leq x)] \\
 &= R_f \int_0^1 A u^\alpha \mathbb{1}(B u^{-\beta} \leq x) du \\
 &= R_f A \int_0^1 u^\alpha \mathbb{1}\left(u \geq \left(\frac{x}{B}\right)^{-\frac{1}{\beta}}\right) du \\
 &= \frac{R_f A}{\alpha + 1} \left(1 - \left(\frac{x}{B}\right)^{-\frac{\alpha+1}{\beta}}\right) \\
 &= 1 - \left(\frac{x}{B}\right)^{-\frac{\alpha+1}{\beta}}.
 \end{aligned}$$

The last line follows from (A.21) below.

- (iii) Routine calculations show that the mean and variance of R_m are given by (provided $\beta < 1/2$)

$$\mathbb{E}[R_m] = \frac{B}{1-\beta} \quad \sigma^2(R_m) = \frac{B^2}{1-2\beta} - \left(\frac{B}{1-\beta}\right)^2. \quad (\text{A.19})$$

Likewise, the distribution of the SDF follows from

$$\mathbb{P}(M \leq x) = \mathbb{P}(A U^\alpha \leq x) = \left(\frac{x}{A}\right)^{\frac{1}{\alpha}}, \quad 0 \leq x \leq A.$$

In this case, M is said to have a Pareto lower tail. The expectation is given by

$$\mathbb{E}[M] = \frac{A}{\alpha + 1}.$$

The constraint $\mathbb{E}[MR_m] = 1$ forces

$$\frac{AB}{\alpha - \beta + 1} = 1. \quad (\text{A.20})$$

In addition from $\mathbb{E}[M] = \frac{1}{R_f}$ it follows that

$$\frac{A}{\alpha + 1} = \frac{1}{R_f}. \quad (\text{A.21})$$

The Sharpe ratio can now be computed from (A.19) and (A.21).

- (iv) It is straightforward to show that the quantiles of a $\mathbf{Par}(C, \zeta)$ distribution are given by

$$Q_\tau = C(1 - \tau)^{-1/\zeta}.$$

It therefore follows that the risk-neutral quantile function is equal to

$$\tilde{Q}_\tau = B(1 - \tau)^{-\frac{\beta}{\alpha+1}}.$$

As a result

$$\begin{aligned} \mathbb{P}(R_m \leq \tilde{Q}_\tau) &= \mathbb{P}\left(R_m \leq B(1 - \tau)^{-\frac{\beta}{\alpha+1}}\right) \\ &= 1 - \left(\frac{B}{B(1 - \tau)^{-\frac{\beta}{\alpha+1}}}\right)^{\frac{1}{\beta}} \\ &= 1 - (1 - \tau)^{\frac{1}{\alpha+1}}. \end{aligned}$$

Hence, the distribution bound evaluates to

$$\frac{1}{R_f} \frac{|\tau - \phi(\tau)|}{\sqrt{\phi(\tau)(1 - \phi(\tau))}} = \frac{A}{1 + \alpha} \frac{\left|\tau - 1 + (1 - \tau)^{\frac{1}{\alpha+1}}\right|}{\sqrt{(1 - (1 - \tau)^{\frac{1}{\alpha+1}})(1 - \tau)^{\frac{1}{\alpha+1}}}.$$

- (v) The HJ bound, as given by the Sharpe ratio in (A.18), goes to 0 as $\beta \nearrow 1/2$ since $\sigma(R_m) \nearrow \infty$. ■

A.9 Derivation of Gâteaux Derivative

In this Section I derive (6.3). For ease of exposition, I drop the time subscripts. For $\lambda \in [0, 1]$, define $\tilde{F}_\lambda := (1 - \lambda)\tilde{F} + \lambda F$. The following (trivial) identity will prove

helpful³²

$$\tau = \tilde{F}_\lambda \tilde{F}_\lambda^{-1}. \quad (\text{A.22})$$

To further simplify notation, write $q(\lambda) := \tilde{F}_\lambda^{-1}$. Then (A.22) becomes

$$\tau = (1 - \lambda)\tilde{F}(q(\lambda)) + \lambda F(q(\lambda)).$$

Applying the implicit function theorem, we obtain

$$q'(\lambda) = -\frac{-\tilde{F}(q(\lambda)) + F(q(\lambda))}{(1 - \lambda)\tilde{f}(q(\lambda)) + \lambda f(q(\lambda))}.$$

Plug in $\lambda = 0$ to get

$$q'(0) = -\frac{-\tilde{F}(q(0)) + F(q(0))}{\tilde{f}(q(0))}. \quad (\text{A.23})$$

Notice that

$$\tilde{F}_\lambda|_{\lambda=0} = \tilde{F} \implies q(\lambda)|_{\lambda=0} = q(0) = \tilde{F}^{-1}. \quad (\text{A.24})$$

Substitute (A.24) into (A.23) to obtain

$$q'(0) = -\frac{-\tilde{F}(\tilde{F}^{-1}) + F(\tilde{F}^{-1})}{\tilde{f}(\tilde{F}^{-1})} = \frac{\tau - F(\tilde{F}^{-1})}{\tilde{f}(\tilde{F}^{-1})}. \quad (\text{A.25})$$

Notice that $q'(0)$ is exactly equal to the Gâteaux derivative from the definition in (6.2), since

$$\frac{\partial}{\partial \lambda} \varphi \left[(1 - \lambda)\tilde{F} + \lambda F \right] \Big|_{\lambda=0} = \frac{\partial}{\partial \lambda} q(\lambda) \Big|_{\lambda=0} = q'(0).$$

A.10 Proof of Proposition 6.4

In the proofs that follow, I repeatedly use Taylor's theorem with integral remainder, which is stated for completeness.

Lemma A.9 (Taylor's theorem). *Let $\zeta^{(3)}(\cdot)$ be absolutely continuous on the closed interval between a and x , then*

$$\zeta(x) = \sum_{k=0}^3 \frac{\zeta^{(k)}(a)}{k!} (x - a)^k + \int_a^x \frac{\zeta^{(4)}(t)}{3!} (x - t)^3 dt.$$

The proof of Proposition 6.4 proceeds in several stages, by first proving an infeasible lower bound on $\tau - F_t(\tilde{Q}_{t,\tau})$, which is later refined into a feasible lower bound under additional assumptions. Before doing so, I collect several results about the SDF in representative agent models.

³²This “equality” may actually only be an inequality for some τ , but this is immaterial to the argument.

Lemma A.10. Assume a representative agent model with SDF given by (6.5), then

$$\tau - F_t(\tilde{Q}_{t,\tau}) = -\frac{\widetilde{\text{COV}}_t \left[\mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right), \zeta(R_{m,t \rightarrow N}) \right]}{\widetilde{\mathbb{E}}_t [\zeta(R_{m,t \rightarrow N})]}, \quad (\text{A.26})$$

where $\zeta(\cdot)$ is defined in (6.6).

Proof. Use the reciprocal of the SDF to pass from physical to risk-neutral measure

$$\begin{aligned} F_t(\tilde{Q}_{t,\tau}) &= \mathbb{E}_t \left[\mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right) \right] = \widetilde{\mathbb{E}}_t \left[\mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right) \frac{\mathbb{E}_t [M_{t \rightarrow N}]}{M_{t \rightarrow N}} \right] \\ &= \widetilde{\text{COV}}_t \left[\mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right), \frac{\mathbb{E}_t [M_{t \rightarrow N}]}{M_{t \rightarrow N}} \right] + \tau. \end{aligned} \quad (\text{A.27})$$

Rearranging the above and using the definition of $\zeta(\cdot)$ in (6.6), as well as (6.5), we obtain (A.26). \blacksquare

Lemma A.11. Under Assumption 6.2,

$$\widetilde{\mathbb{E}}_t [\zeta(R_{m,t \rightarrow N})] \leq \sum_{k=0}^3 \theta_k \widetilde{\mathbb{M}}_{t \rightarrow N}^{(k)} = 1 + \sum_{k=1}^3 \theta_k \widetilde{\mathbb{M}}_{t \rightarrow N}^{(k)},$$

where $\zeta(x)$ is the IMRS defined in (6.6).

Proof. In the integral of Lemma A.9, substitute $s = (t - a)/(x - a)$ to get

$$\begin{aligned} \zeta(x) &= \sum_{k=0}^3 \frac{\zeta^{(k)}(a)}{k!} (x - a)^k + (x - a)^4 \int_0^1 \frac{\zeta^{(4)}(a + s(x - a))}{3!} (1 - s)^3 ds \\ &\leq \sum_{k=0}^3 \frac{\zeta^{(k)}(a)}{k!} (x - a)^k, \end{aligned}$$

since $\zeta^{(4)}(x) < 0$ by Assumption 6.2(ii). Using this result with $a = R_{f,t \rightarrow N}$ and taking expectations, we obtain

$$\widetilde{\mathbb{E}}_t [\zeta(R_{m,t \rightarrow N})] \leq \sum_{k=0}^3 \theta_k \widetilde{\mathbb{M}}_{t \rightarrow N}^{(k)}. \quad \blacksquare$$

Under Assumption 6.2, the difference between the physical and risk-neutral distribution in the left-tail can be bounded as follows.

Theorem A.12 (Infeasible Lower Bound). *Let Assumption 6.2 hold and assume that the risk-neutral CDF is absolutely continuous with respect to Lebesgue measure. Define τ^* so that $G(\tilde{Q}_{t,\tau^*}) = \widetilde{\mathbb{E}}_t (G(R_{m,t \rightarrow N}))$, where*

$$G(R_{m,t \rightarrow N}) := \int_{R_{f,t \rightarrow N}}^{R_{m,t \rightarrow N}} \zeta^{(4)}(t) (R_{m,t \rightarrow N} - t)^3 dt.$$

Then for all $\tau \leq \tau^*$,

$$\tau - F_t(\tilde{Q}_{t,\tau}) \geq \frac{\sum_{k=1}^3 \theta_k \left(\tau \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)} - \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)}[\tilde{Q}_{t,\tau}] \right)}{1 + \sum_{k=1}^3 \theta_k \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)}}, \quad (\text{A.28})$$

where $\tilde{\mathbb{M}}_{t \rightarrow N}^{(k)}, \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)}[\tilde{Q}_{t,\tau}]$ are defined in (6.7).

Proof of Theorem A.12. By Taylor's theorem,

$$\begin{aligned} & -\widetilde{\text{COV}}_t \left[\mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right), \zeta(R_{m,t \rightarrow N}) \right] = \sum_{k=1}^3 \theta_k \left(\tau \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)} - \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)}[\tilde{Q}_{t,\tau}] \right) \\ & -\widetilde{\text{COV}}_t \left[\mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right), \frac{1}{3!} \int_{R_{f,t \rightarrow N}}^{R_{m,t \rightarrow N}} \zeta^{(4)}(t) (R_{m,t \rightarrow N} - t)^3 dt \right] \\ & \geq \sum_{k=1}^3 \theta_k \left(\tau \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)} - \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)}[\tilde{Q}_{t,\tau}] \right). \end{aligned} \quad (\text{A.29})$$

The last line follows from Lemma A.13 below. Hence,

$$\begin{aligned} \tau - F_t(\tilde{Q}_{t,\tau}) &= -\frac{\widetilde{\text{COV}}_t \left[\mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right), \zeta(R_{m,t \rightarrow N}) \right]}{\tilde{\mathbb{E}}_t [\zeta(R_{m,t \rightarrow N})]} \\ &\geq \frac{\sum_{k=1}^3 \theta_k \left(\tau \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)} - \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)}[\tilde{Q}_{t,\tau}] \right)}{1 + \sum_{k=1}^3 \theta_k \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)}}, \end{aligned}$$

where the first identity follows from Lemma A.10 and the inequality follows from (A.29) and Lemma A.11. \blacksquare

Remark 6. The condition that $\tau \leq \tau^*$ is sufficient but not necessary, as the proof of Theorem A.12 shows. Furthermore, the proof also shows that $\tau^* > 0$ exists regardless of the utility function. In practice, however, τ^* is unknown since $G(\cdot)$ depends on the unknown utility function of the representative agent. Appendix D.5 shows that $\tau^* \approx 0.5$ in the data for CRRA utility and different levels of risk aversion. In light of this result, it seems that $\tau \in \{0.05, 0.1, 0.2\}$ is sufficiently conservative for the lower bound to hold, and I use these values in the empirical application in Section 6.4.

Lemma A.13. *Suppose that Assumption 6.2 holds. In addition, define τ^* so that $G(\tilde{Q}_{t,\tau^*}) = \tilde{\mathbb{E}}_t (G(R_{m,t \rightarrow N}))$, where*

$$G(R_{m,t \rightarrow N}) := \int_{R_{f,t \rightarrow N}}^{R_{m,t \rightarrow N}} \zeta^{(4)}(t) (R_{m,t \rightarrow N} - t)^3 dt.$$

Then for all $\tau \leq \tau^*$,

$$\widetilde{\text{COV}}_t \left[\mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right), \int_{R_{f,t \rightarrow N}}^{R_{m,t \rightarrow N}} \zeta^{(4)}(t) (R_{m,t \rightarrow N} - t)^3 dt \right] \leq 0. \quad (\text{A.30})$$

Proof. If $\zeta^{(4)} \equiv 0$, then (A.30) trivially holds. Hence, assume that $\zeta^{(4)}$ is not identically equal to zero. First we show that $G(R_{m,t \rightarrow N})$ is increasing on $(0, R_{f,t \rightarrow N})$, since by Leibniz' rule

$$G'(R_{m,t \rightarrow N}) = -3 \int_{R_{f,t \rightarrow N}}^{R_{m,t \rightarrow N}} \zeta^{(4)}(t) (R_{m,t \rightarrow N} - t)^2 dt \geq 0.$$

The inequality follows since $\zeta^{(4)}(t) < 0$ by Assumption 6.2(ii). Temporarily write $K = \tilde{Q}_{t,\tau}$ to ease notation and consider

$$\Gamma(K) = \widetilde{\text{COV}}_t \left[\mathbb{1} (R_{m,t \rightarrow N} \leq K), \int_{R_{f,t \rightarrow N}}^{R_{m,t \rightarrow N}} \zeta^{(4)}(t) (R_{m,t \rightarrow N} - t)^3 dt \right].$$

By Leibniz' rule again, we get

$$\Gamma'(K) = \tilde{f}_t(K) \left(G(K) - \tilde{\mathbb{E}}_t (G(R_{m,t \rightarrow N})) \right).$$

Since $G(R_{f,t \rightarrow N}) = 0$, $G(R_{m,t \rightarrow N}) \leq 0$ and $G(R_{m,t \rightarrow N})$ is increasing on $(0, R_{f,t \rightarrow N})$, we know that $\Gamma'(K) \leq 0$ for all $K \leq K^* < R_{f,t \rightarrow N}$, where K^* is defined such that $G(K^*) = \tilde{\mathbb{E}}_t (G(R_{m,t \rightarrow N}))$. To complete the proof, define τ^* so that it satisfies $\tilde{Q}_{\tau^*} = K^*$. ■

Remark 7. The bound in (A.28) is infeasible since $\{\theta_k\}_{k=1}^3$ is unknown.³³ However, Chabi-Yo and Loudis (2020) show that these unknowns relate to the coefficient of relative risk aversion, relative prudence and relative temperance of the representative agent. Based on this observation and using results from the expected utility literature (Eeckhoudt and Schlesinger, 2006), the authors propose an additional restriction on θ_k that allows me to prove the feasible lower bound in Proposition 6.4.

Proof of Proposition 6.4. Using Assumption 6.3(i) and 6.3(ii), we get $\theta_2 \tilde{\mathbb{M}}_{t \rightarrow N}^{(2)} \leq -1/R_{f,t \rightarrow N}^2 \tilde{\mathbb{M}}_{t \rightarrow N}^{(2)}$ and $\theta_3 \tilde{\mathbb{M}}_{t \rightarrow N}^{(3)} \leq 1/R_{f,t \rightarrow N}^3 \tilde{\mathbb{M}}_{t \rightarrow N}^{(3)}$, from which it follows that

$$1 + \sum_{k=1}^3 \theta_k \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)} \leq 1 - \frac{1}{R_{f,t \rightarrow N}^2} \tilde{\mathbb{M}}_{t \rightarrow N}^{(2)} + \frac{1}{R_{f,t \rightarrow N}^3} \tilde{\mathbb{M}}_{t \rightarrow N}^{(3)}. \quad (\text{A.31})$$

³³In Appendix E, I use comparative statics for common utility functions to analyze the tail difference between the physical and risk-neutral distribution.

Second, recall that for $K > 0$

$$\tilde{F}_t(K)\tilde{\mathbb{M}}_{t \rightarrow N}^{(k)} - \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)}[K] = -\widetilde{\mathbb{COV}}_t \left[\mathbb{1}(R_{m,t \rightarrow N} \leq K), (R_{m,t \rightarrow N} - R_{f,t \rightarrow N})^k \right].$$

If $k = 1, 3$, then Chebyshev's sum inequality [E.1](#) implies that

$$\Gamma(K) := \widetilde{\mathbb{COV}}_t \left[\mathbb{1}(R_{m,t \rightarrow N} \leq K), (R_{m,t \rightarrow N} - R_{f,t \rightarrow N})^k \right] \leq 0.$$

Hence under Assumption [6.3\(i\)](#),

$$\theta_k \left(\tilde{F}_t(K)\tilde{\mathbb{M}}_{t \rightarrow N}^{(k)} - \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)}[K] \right) \geq \frac{1}{R_{f,t \rightarrow N}^k} \left(\tilde{F}_t(K)\tilde{\mathbb{M}}_{t \rightarrow N}^{(k)} - \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)}[K] \right) \quad \text{for } k = 1, 3. \quad (\text{A.32})$$

If $k = 2$, we obtain from Leibniz' rule

$$\Gamma'(K) = \tilde{f}_t(K) \left[(K - R_{f,t \rightarrow N})^2 - \widetilde{\mathbb{VAR}}_t(R_{m,t \rightarrow N}) \right]. \quad (\text{A.33})$$

It follows that [\(A.33\)](#) is positive if $K \leq R_{f,t \rightarrow N} - \sqrt{\widetilde{\mathbb{VAR}}_t(R_{m,t \rightarrow N})} =: K^{**}$. Combining [\(A.32\)](#) and [\(A.33\)](#), we get for $K \leq K^{**}$

$$\theta_k \left(\tilde{F}_t(K)\tilde{\mathbb{M}}_{t \rightarrow N}^{(k)} - \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)}[K] \right) \geq \frac{(-1)^{k+1}}{R_{f,t \rightarrow N}^k} \left(\tilde{F}_t(K)\tilde{\mathbb{M}}_{t \rightarrow N}^{(k)} - \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)}[K] \right). \quad (\text{A.34})$$

Collecting the results from [\(A.31\)](#) and [\(A.34\)](#) and using the general upper bound [\(A.28\)](#) from Theorem [A.12](#), it follows that

$$\begin{aligned} \tau - F_t(\tilde{Q}_{t,\tau}) &\stackrel{(\text{A.28})}{\geq} \frac{\sum_{k=1}^3 \theta_k \left(\tau \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)} - \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)}[\tilde{Q}_{t,\tau}] \right)}{1 + \sum_{k=1}^3 \theta_k \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)}} \\ &\geq \frac{\sum_{k=1}^3 \frac{(-1)^{k-1}}{R_{f,t \rightarrow N}^k} \left(\tau \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)} - \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)}[\tilde{Q}_{t,\tau}] \right)}{1 + \sum_{k=1}^3 \frac{(-1)^{k-1}}{R_{f,t \rightarrow N}^k} \tilde{\mathbb{M}}_{t \rightarrow N}^{(k)}}, \end{aligned}$$

for all τ such that $\tilde{Q}_{t,\tau} \leq \min(K^*, K^{**})$, where K^* is defined in Theorem [A.12](#). \blacksquare

Remark 8. The bound only holds for quantiles far enough in the left-tail. Compared to Theorem [A.12](#), the additional condition needed for the bound to hold is that $\tilde{Q}_{t,\tau} \leq R_{f,t \rightarrow N} - \sqrt{\widetilde{\mathbb{VAR}}_t(R_{m,t \rightarrow N})}$, which covers a wide range of quantiles in the left-tail, since in the data $\sqrt{\widetilde{\mathbb{VAR}}_t(R_{m,t \rightarrow N})}$ is in the order of 10^{-3} for 90-day returns, whereas the risk-free rate is typically around 1.^{[34](#)}

³⁴At the 30- and 60-day horizon, the risk-neutral standard deviation is even smaller.

A.11 Formulas for market moments

This Section presents formulas for the (un)truncated risk-neutral moments of the excess market return. I use a slight abuse of notation and write $\tilde{Q}(\tau) := \tilde{Q}_\tau(R_{m,t \rightarrow N})$, to emphasize that the integrals below are taken with respect to τ .

Proposition A.14. *Any risk-neutral moment can be computed from the risk-neutral quantile function, since*

$$\tilde{\mathbb{E}}_t [(R_{m,t \rightarrow N} - R_{f,t \rightarrow N})^n] = \int_0^1 [\tilde{Q}_\tau(R_{m,t \rightarrow N} - R_{f,t \rightarrow N})]^n d\tau = \int_0^1 [\tilde{Q}(\tau) - R_{f,t \rightarrow N}]^n d\tau. \quad (\text{A.35})$$

Moreover, any truncated risk-neutral moment can be calculated by

$$\tilde{\mathbb{E}}_t [(R_{m,t \rightarrow N} - R_{f,t \rightarrow N})^n \mathbb{1}(R_{m,t \rightarrow N} \leq k_0)] = \int_0^{\tilde{F}_t(k_0)} [\tilde{Q}(\tau) - R_{f,t \rightarrow N}]^n d\tau.$$

Proof. For any random variable X and integer n such that the n -th moment exists, we have

$$\mathbb{E}[X^n] = \int_0^1 [Q_X(\tau)]^n d\tau.$$

This follows straightforward from the substitution $x = Q(\tau)$. Now use that for any constant $a \in \mathbb{R}$, $Q_{X-a}(\tau) = Q_X(\tau) - a$ to derive (A.35). The truncated formula follows similarly. \blacksquare

Remark 9. Frequently I use $k_0 = \tilde{Q}_\tau$, in which case the truncated moment formula reduces to

$$\tilde{\mathbb{E}}_t [(R_{m,t \rightarrow N} - R_{f,t \rightarrow N})^n \mathbb{1}(R_{m,t \rightarrow N} \leq \tilde{Q}_\tau)] = \int_0^\tau [\tilde{Q}(p) - R_{f,t \rightarrow N}]^n dp.$$

B Risk-Neutral Quantile Regression: Robustness and Departure from Conditional Lognormality

B.1 Linear versus Non-linear Model: Out-of-Sample Forecasting Accuracy

This section explores alternative specifications to the linear quantile model presented in (2.3), focusing only on 30-day returns. The findings for longer time horizons are very similar and omitted for parsimony. Specifically, I consider the addition of higher-order terms to the linear model, such as:

$$Q_{t,\tau} = \beta_0(\tau) + \beta_1(\tau)\tilde{Q}_{t,\tau} + \beta_2(\tau)\tilde{Q}_{t,\tau}^2. \quad (\text{B.1})$$

To evaluate the performance of the non-linear model in (B.1) vs. the linear model in (2.3), I recursively estimate the model parameters based on an expanding window, starting at January 2, 2003. The first sub-sample ends at August 15, 2012 and I increase the sample size on a monthly basis. For each sub-sample, I calculate the out-of-sample forecasting accuracy using the formula:

$$\frac{1}{\#t} \sum_t \rho_\tau(R_{m,t \rightarrow N} - \hat{Q}_{t,\tau}), \quad (\text{B.2})$$

where $\hat{Q}_{t,\tau}$ is the predicted physical quantile based on the parameters estimated from the sub-sample. The summation includes all dates that are at least one month ahead of the end of the sub-sample period.

Figure B3 shows the out-of-sample loss at various percentiles. In most cases, the linear model outperforms the quadratic model, with some exceptions observed at the 95th percentile during specific periods. These results continue to hold when adding other non-linear terms, such as cubic, exponential or logarithmic factors. Additionally, I find that the risk-neutral quantile function exhibits a high correlation with higher-order terms. Consequently, the non-linear model tends to produce quantile forecasts that closely resemble those generated by the linear model.

B.2 Additional Evidence Against the Lognormal Assumption

Table 2 already indicates evidence against the lognormal model since the QR estimates in the left- and right-tail are rather different, in contradiction with (4.4). To further assess the implications of the lognormal model, I analyze the accuracy of the physical quantile forecast in (4.5) out-of-sample. Specifically, I use QR based on the first t_0 observations to estimate the model

$$Q_{t,\tau}(R_{m,t \rightarrow N}) = \hat{\beta}_{0,t_0}(\tau) + \hat{\beta}_{1,t_0}(\tau) \tilde{Q}_{t,\tau}, \quad (\text{B.3})$$

where the t_0 -subscript in β_{\cdot,t_0} refers to the fact that the coefficients are estimated using observations up to time t_0 . Using an expanding window to estimate β_{\cdot,t_0} , the model produces dynamic quantile forecasts of the form

$$\hat{Q}_{t,\tau}^{\text{logn}} = \hat{\beta}_{0,t}(\tau) + \hat{\beta}_{1,t}(\tau) \tilde{Q}_{t,\tau}. \quad (\text{B.4})$$

In the lognormal case, Proposition 4.1(ii) suggests that $Q_{t,\tau}(R_{m,t \rightarrow N}) \approx \hat{Q}_{t,\tau}^{\text{logn}}$. This approximation can be tested using the joint restriction

$$H_0 : [\beta_0(\tau), \beta_1(\tau)] = [0, 1],$$




Figure B3: **Out-of-sample quantile forecasting loss.** These figures show the out-of-sample loss of forecasting the physical quantile function, based on an expanding window. The loss at different percentiles is calculated by (B.2).

in the quantile regression

$$\min_{\beta_0, \beta_1 \in \mathbb{R}} \sum_t \rho_\tau \left(R_{m,t \rightarrow N} - \beta_0 - \beta_1 \hat{Q}_{t,\tau}^{\text{logn}} \right).$$

The results are summarized in Table B1 and show that the point estimates are quite far from the $[0, 1]$ benchmark. The Wald test on the joint restriction tends to reject H_0 far enough in the tail, but for $\tau = 0.2$ the null hypothesis is never rejected due to the large standard errors. Additionally, the $R^1(\tau)$ statistic shows that the explanatory power is low relative to Table 2, even though the sample sizes are different. Hence, the results are incompatible with (4.4) and (4.5) and provide evidence against the conditional lognormal assumption, which is in line with evidence from the literature (see e.g. Martin (2017, Result 4)).

Table B1: **Expanding quantile prediction with risk-neutral quantile**

Horizon	τ	$\hat{\beta}_0(\tau)$	$\hat{\beta}_1(\tau)$	Wald test (p -value)	$R^1(\tau)[\%]$	Obs
<u>30 days</u>	0.05	0.54 (0.185)	0.42 (0.193)	0.00	4.36	3804
	0.1	0.59 (0.205)	0.39 (0.212)	0.01	2.39	
	0.2	0.83 (0.332)	0.15 (0.338)	0.04	0.28	
<u>60 days</u>	0.05	0.55 (0.310)	0.39 (0.329)	0.06	1.84	3753
	0.1	0.80 (0.339)	0.16 (0.352)	0.03	0.22	
	0.2	0.87 (0.416)	0.11 (0.425)	0.10	0.25	
<u>90 days</u>	0.05	0.78 (0.335)	0.11 (0.358)	0.01	0.88	3702
	0.1	0.74 (0.376)	0.21 (0.395)	0.08	1.34	
	0.2	0.73 (0.481)	0.26 (0.491)	0.31	0.52	

Note: This table reports the QR estimates of (B.4) using an expanding window based on an initial 500 observations. The sample period is 2003-2021. *Wald test* denotes the p -value of the joint restriction $[\beta_0(\tau), \beta_1(\tau)] = [0, 1]$. Standard errors are reported in parentheses and calculated using the SETBB with a block length equal to the prediction horizon. $R^1(\tau)$ denotes the goodness of fit measure (2.5).

C Estimating the Risk-Neutral Quantile Function

C.1 Data Description

To estimate the risk-neutral quantile function at each time point, I use daily option prices from OptionMetrics, covering the period from January 1, 1996, to December 31, 2021. These options include European Put and Call options on the S&P500 index. The option contracts provide data on the highest closing bid, lowest closing ask price, and the price of the forward contract on the underlying security. To approximate the unobserved option price, I use the midpoint between the bid and ask prices. Additionally, I gather daily risk-free rate data from Kenneth French's website.³⁵ Finally, I obtain stock price data on the closing price of the S&P500 from WRDS.

I implement an additional data cleaning procedure for the option data before estimating the martingale measure. Firstly, I exclude all observations where the highest closing bid price is zero. Additionally, I remove option prices that violate no-

³⁵See http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html#Research

arbitrage bounds. Subsequently, I filter out option prices with maturities less than 7 days or greater than 500 days. Following this cleaning process, I retain 23,264,113 option-day observations.

For the quantile regression application, I exclude all observations before 2003. During the period from 1996 to 2003, there are many days with insufficient option data to estimate $\tilde{Q}_{t,\tau}$ at the 30-, 60-, and 90-day horizons. I also discard days in the post-2003 period when I cannot estimate the risk-neutral quantile, although this is a rare occurrence, accounting for approximately 2% of the total days. Most of these instances are concentrated at the beginning of the sample period.

C.2 Estimation Procedure

There is a substantial literature on how to extract the martingale measure from option prices. I use the **RND Fitting Tool** application on MATLAB, which is developed by [Barletta and Santucci de Magistris \(2018\)](#).³⁶ The tool is based on the orthogonal polynomial expansion of [Filipović et al. \(2013\)](#). In short, the idea is to approximate the conditional risk-neutral density function by an expansion of the form

$$\tilde{f}_t(x) \approx \phi(x) \left[1 + \sum_{k=1}^K \sum_{i=0}^k c_k w_{i,k} x^k \right],$$

where $\phi(x)$ is an arbitrary density and the polynomial term serves to tilt the density function towards the risk-neutral distribution. Further details about the estimation of the coefficients $w_{i,k}$ and c_k can be found in [Filipović et al. \(2013\)](#).

For my purpose, I need to choose the kernel function $\phi(\cdot)$, the estimation method for c_k and the degree of the expansion K . I follow the recommendation of [Barletta and Santucci de Magistris \(2018\)](#) and use the double beta distribution for the kernel and principal component analysis to estimate c_k . This is the most robust method for S&P500 options. To avoid overfitting, I use $K = 3$ if the number of option data is less than 70, $K = 6$ if the number is less than 100 and $K = 8$ otherwise. This choice renders a good approximation for most time periods.

I interpolate the estimated risk-neutral densities for a given time horizon. Occasionally, there are no two interpolation points. In such cases, I drop the observations to avoid negative density estimates due to extrapolation. Since the RND Fitting Tool is designed for an equal number of put and call options, I use Put-Call parity to convert in-the-money call prices to put prices and vice versa. Subsequently, I use

³⁶The application can be downloaded from the author's GITHUB page: <https://github.com/abarletta/rndfittool>

Black-Scholes implied volatilities to interpolate the Call-Put option price curve near the forward price. This transformation ensures that the risk-neutral density does not have a discontinuity for strike prices that are close to being at-the-money (Figlewski, 2010). Finally, I integrate the density function and take the inverse to obtain the risk-neutral quantile function:

$$\tilde{Q}_{t,\tau} := \inf \left\{ x \in \mathbb{R} : \tau \leq \tilde{F}_t(x) \right\}, \quad \text{where } \tilde{F}_t(x) = \int_0^x \tilde{f}_t(y) dy.$$

D Verifying Assumption 6.2(ii) in Representative Agent Models

The proof of Theorem A.12 relies on Assumption 6.2(ii). This section derives parameter restrictions for common utility functions that are needed so that Assumption 6.2(ii) is satisfied. Most of these restrictions resemble those of Chabi-Yo and Loudis (2020). I also illustrate the lower bound with actual data assuming CRRA utility.

D.1 Log utility

In this case $u(x) = \log x$. It follows that $\zeta(x) = x/R_{f,t \rightarrow N}$. Clearly $\zeta^{(4)}(x) = 0$ and Assumption 6.2 holds.

D.2 CRRA utility

More generally, consider $u(x) = \frac{x^{1-\gamma}}{1-\gamma}$ for $\gamma \geq 0$. It follows that $\zeta(x) = (\frac{x}{R_{f,t \rightarrow N}})^\gamma$ and hence

$$\zeta^{(4)}(x) = \frac{1}{R_{f,t \rightarrow N}^\gamma} \gamma(\gamma-1)(\gamma-2)(\gamma-3)x^{\gamma-4}.$$

Part (ii) of Assumption 6.2 holds if $\gamma \in [0, 1]$, but also if $\gamma \in [2, 3]$. Notice that the additional restrictions in the feasible lower bound in Proposition 6.4 cannot be accommodated by this model. To see this, observe that $\theta_2 \leq -1/R_{f,t \rightarrow N}^2$ implies that $\gamma(\gamma-1)/2 \leq -1/R_{f,t \rightarrow N}^2$, which cannot hold for any reasonable interest rate. This failure illustrates that a representative agent model with CRRA utility is misspecified in that it cannot produce a sizable risk-premium on skewness.³⁷

D.3 CARA utility

In this case, $u(x) = 1 - e^{-\gamma x}$ and $\zeta(x) = e^{\gamma^*(x - R_{f,t \rightarrow N})}$, where $\gamma^* = W_t \gamma$. Since $\zeta^{(4)} > 0$, Assumption 6.2 does not hold.

³⁷See in particular Chabi-Yo and Loudis (2020, Equation (A.5)), which shows that θ_2 is related to the risk-premium on market skewness.

D.4 HARA utility

The utility function is given by $u(x) = \frac{1-\gamma}{\gamma} \left(\frac{ax}{1-\gamma} + b \right)^\gamma$, where $a > 0$ and $\frac{ax}{1-\gamma} + b > 0$. Successive differentiation renders

$$\zeta^{(4)}(x) = \frac{-\gamma(\gamma+1)(\gamma+2)(aW_t)^4 \left(\frac{aW_t x}{1-\gamma} + b \right)^{-\gamma-3} \left(\frac{aW_t R_{f,t \rightarrow N}}{1-\gamma} + b \right)^{\gamma-1}}{(1-\gamma)^3}.$$

We see that $\gamma \in [0, 1)$ is a sufficient condition for $\zeta^{(4)}(x) \leq 0$.

D.5 Lower Bound in the Data for CRRA utility

Figure D4 illustrates the infeasible lower bound as well as the quantile approximation for CRRA utility with different levels of risk aversion. The risk-neutral distribution is obtained from option data over a 90-day horizon on October 28, 2015. Panels D4a and D4c show the infeasible lower bound from Theorem A.12 when risk aversion is 2.2 and 2.9 respectively. Consistent with the theorem, the infeasible lower bound is below $\tau - F_t(\tilde{Q}_{t,\tau})$ in the left-tail, and seems to hold for a large range of τ 's, in particular for all $\tau \leq 0.5$. The right panels show the quantile approximation (6.9) based on the infeasible lower bound. We see that the risk-adjusted quantile approximation comes much closer to the physical quantile relative to the risk-neutral quantile function.

E Disaster Probability in Representative Agent Models

In this section, I derive results regarding conditional tail probabilities in representative agent models. I demonstrate how these probabilities can be computed using common utility functions and analyze their sensitivity to changes in underlying parameters (comparative statics). These results do not hinge on specific assumptions about the market return distribution and extend existing findings in the literature, which often rely on log-normality assumptions.

E.1 Log Utility

Chabi-Yo and Loudis (2020, Remark 1) show that their bounds on the equity premium equal the bounds of Martin (2017) when the representative agent has log preferences. Here, I derive the analogous result for the subjective crash probability of a log investor reported by Martin (2017, Result 2). In our notation, Martin (2017)




Figure D4: **Lower bound with CRRA utility for 90-day returns.** This figure shows the lower bound on $\tau - F_t(\tilde{Q}_{t,\tau})$ as well as the quantile approximation $Q_{t,\tau} \approx \tilde{Q}_{t,\tau} + \text{LB}_{t,\tau}$ in a representative agent model with CRRA utility function, $u(x) = x^{1-\gamma}/(1-\gamma)$, for $\gamma \in \{2.2, 2.9\}$. The left panels show the infeasible lower bound $\text{CLB}_{t,\tau}$, and the true risk-adjustment, $\tau - F_t(\tilde{Q}_{t,\tau})$. The right panels show the physical, risk-neutral and risk-adjusted quantile functions. The risk-adjusted quantile function uses the infeasible lower bound. The risk-neutral distribution is coming from option data on the S&P500 on October 28, 2015 with a maturity of 90 days.

shows that

$$\mathbb{P}_t(R_{m,t \rightarrow N} < \alpha) = \alpha \left[\text{Put}'_t(\alpha S_t) - \frac{\text{Put}_t(\alpha S_t)}{\alpha S_t} \right], \quad (\text{E.1})$$

where Put'_t is the derivative of the put option price curve seen as a function of the strike. Under log preferences and using (A.27), it follows that

$$\begin{aligned}\mathbb{P}_t(R_{m,t \rightarrow N} < \tilde{Q}_{t,\tau}) &= \tau + \frac{1}{R_{f,t \rightarrow N}} \widetilde{\text{COV}}_t \left[\mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right), R_{m,t \rightarrow N} \right] \\ &= \tau + \frac{1}{R_{f,t \rightarrow N}} \left(\tilde{\mathbb{E}}_t \left[\mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right) R_{m,t \rightarrow N} \right] - \tilde{\mathbb{E}}_t(R_{m,t \rightarrow N}) \tilde{\mathbb{E}}_t \left(\mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right) \right) \right) \\ &= \frac{1}{R_{f,t \rightarrow N}} \tilde{\mathbb{E}}_t \left[\mathbb{1} \left(R_{m,t \rightarrow N} \leq \tilde{Q}_{t,\tau} \right) R_{m,t \rightarrow N} \right].\end{aligned}\tag{E.2}$$

The result now follows upon substituting $\tilde{Q}_\tau = \alpha$, since Martin (2017) shows that (E.2) equals the right hand side of (E.1).

E.2 CRRA Utility

I now consider the case in which the representative agent has constant relative risk aversion (CRRA) utility, $u(x) = x^{1-\gamma}/(1-\gamma)$, where γ is the relative risk aversion parameter. First, I show that the excess market return is non-decreasing in γ *regardless* of the distribution of the market return.³⁸ Next, I extend the argument to show that the difference between the physical and risk-neutral measures increases at every point within their support. The proofs rely on the following lemma, which is a special case of the FKG inequality (Hsu and Varadhan, 1999, Theorem 1.3).

Lemma E.1 (Chebyshev sum inequality). *Let X be a random variable and let g, h both be non-increasing or non-decreasing. Then,*

$$\mathbb{E}(g(X)h(X)) \geq \mathbb{E}(g(X)) \mathbb{E}(h(X)).$$

The inequality is reversed if one is non-increasing and the other is non-decreasing.

Proof. Let X_1, X_2 be IID copies of X and assume that g, h are non-decreasing. It follows that

$$(g(X_1) - g(X_2))(h(X_1) - h(X_2)) \geq 0.\tag{E.3}$$

Taking expectations on both sides completes the proof. The same proof goes through if g, h are non-increasing. If one is non-increasing and the other is non-decreasing, the inequality in (E.3) is reversed. ■

Proposition E.2. *Assume that a representative investor has CRRA utility, with $\gamma \geq 0$ and $\mathbb{E}_t \left[R_{m,t \rightarrow N}^{\gamma+1} \log R_{m,t \rightarrow N} \right] < \infty$. Then, $\mathbb{E}_t[R_{m,t \rightarrow N}] - R_{f,t \rightarrow N}$, is non-decreasing in γ .*

Remark 10. I suppress the dependence of the physical expectation on γ in the notation for readability.

³⁸Cochrane (2005) derives this result when the distribution is lognormal.

Proof. According to Chabi-Yo and Loudis (2020, Equation (53)), we have

$$\mathbb{E}_t[R_{m,t \rightarrow N}] - R_{f,t \rightarrow N} = \frac{\tilde{\mathbb{E}}_t[R_{m,t \rightarrow N}^{\gamma+1}]}{\tilde{\mathbb{E}}_t[R_{m,t \rightarrow N}^\gamma]} - R_{f,t \rightarrow N} =: g(\gamma).$$

It is enough to show that $g'(\gamma) \geq 0$ for $\gamma \geq 0$. Taking first order conditions, we need to show that

$$\tilde{\mathbb{E}}_t[R_{m,t \rightarrow N}^{\gamma+1} \log R_{m,t \rightarrow N}] \tilde{\mathbb{E}}_t[R_{m,t \rightarrow N}^\gamma] \geq \tilde{\mathbb{E}}_t[R_{m,t \rightarrow N}^{\gamma+1}] \tilde{\mathbb{E}}_t[R_{m,t \rightarrow N}^\gamma \log R_{m,t \rightarrow N}]. \quad (\text{E.4})$$

Introduce another probability measure \mathbb{P}^* , defined by

$$\mathbb{E}_t^*[Z] := \frac{\tilde{\mathbb{E}}_t[Z R_{m,t \rightarrow N}^\gamma]}{\tilde{\mathbb{E}}_t[R_{m,t \rightarrow N}^\gamma]}. \quad (\text{E.5})$$

We can rewrite (E.4) into

$$\mathbb{E}_t^*[R_{m,t \rightarrow N}^\gamma \log R_{m,t \rightarrow N}] \geq \mathbb{E}_t^*[R_{m,t \rightarrow N}^\gamma] \mathbb{E}_t^*[\log R_{m,t \rightarrow N}]. \quad (\text{E.6})$$

Inequality (E.6) now follows from Lemma E.1. ■

I mimic the steps above to show that the physical distribution differs more from the risk-neutral distribution at every point in the support, whenever risk aversion is increasing. As before, the dependence of the physical measure on γ is omitted.

Proposition E.3. *Assume that a representative investor has CRRA utility, with $\gamma \geq 0$ and $\mathbb{E}_t[R_{m,t \rightarrow N}^\gamma \log R_{m,t \rightarrow N}] < \infty$, then $F_t(x)$ is non-increasing in γ . In particular, $\tau - F_t(\tilde{Q}_{t,\tau})$ is non-decreasing in γ .*

Proof. I start from the relation

$$F_t(x) = \tilde{\mathbb{E}}_t \left[\frac{R_{m,t \rightarrow N}^\gamma}{\tilde{\mathbb{E}}_t[R_{m,t \rightarrow N}^\gamma]} \mathbb{1}(R_{m,t \rightarrow N} \leq x) \right].$$

From first order conditions, we need to show that

$$\begin{aligned} \tilde{\mathbb{E}}_t \left[\log(R_{m,t \rightarrow N}) \mathbb{1}(R_{m,t \rightarrow N} \leq x) R_{m,t \rightarrow N}^\gamma \right] \tilde{\mathbb{E}}_t \left[R_{m,t \rightarrow N}^\gamma \right] &\leq \\ \tilde{\mathbb{E}}_t \left[R_{m,t \rightarrow N}^\gamma \mathbb{1}(R_{m,t \rightarrow N} \leq x) \right] \tilde{\mathbb{E}}_t \left[\log(R_{m,t \rightarrow N}) R_{m,t \rightarrow N}^\gamma \right]. \end{aligned}$$

Using the same change of measure as in (E.5), we obtain the equivalent statement

$$\mathbb{E}_t^*[\log(R_{m,t \rightarrow N}) \mathbb{1}(R_{m,t \rightarrow N} \leq x)] \leq \mathbb{E}_t^*[\mathbb{1}(R_{m,t \rightarrow N} \leq x)] \mathbb{E}_t^*[\log R_{m,t \rightarrow N}].$$

This inequality holds, since $\log(y)$ and $\mathbb{1}(y \leq x)$ are respectively increasing and non-increasing in y , hence the result follows from Lemma E.1. Using the substitution $x \rightarrow \tilde{Q}_{t,\tau}$, it follows that $\tau - F_t(\tilde{Q}_{t,\tau})$, is non-decreasing in γ . ■

E.3 Exponential utility

Here, I assume that the representative agent has exponential utility, $u(x) = 1 - e^{-\gamma^* x}$, where γ^* is the absolute risk aversion. According to Chabi-Yo and Loudis (2020, Equation (55)), the following expression for the equity premium obtains

$$\mathbb{E}_t[R_{m,t \rightarrow N}] - R_{f,t \rightarrow N} = \frac{\tilde{\mathbb{E}}_t[R_{m,t \rightarrow N} e^{\gamma R_{m,t \rightarrow N}}]}{\tilde{\mathbb{E}}_t[e^{\gamma R_{m,t \rightarrow N}}]} - R_{f,t \rightarrow N},$$

where $\gamma = \gamma^* W_t$ is relative risk aversion and W_t represents the agent's wealth at time t . Since there is a one-to-one relation between γ and γ^* , it follows from the results in Section E.2 that the equity premium is increasing in γ^* , and so is the distance between the physical and risk-neutral distribution, as measured by: $\tau - F_t(\tilde{Q}_{t,\tau})$.

F Lower Bound in the Data and Robustness

F.1 Lower Bound in the Data

In the empirical application, I calculate the lower bound, $LB_{t,\tau} = CLB_{t,\tau} / \tilde{f}_t(\tilde{Q}_{t,\tau})$, for 30-, 60-, and 90-day returns. Summary statistics of $LB_{t,\tau}$ are presented in Table F2. The lower bound is right-skewed and is most significant at the 5th and 10th percentile. Moreover, over the 30-day horizon, it can reach as high as 25% and maintains an average of approximately 1% in the far left-tail.

F.2 Robustness of the Lower Bound and Risk-neutral Quantile

The lower bound, $LB_{t,\tau}$, tries to capture the difference between the physical and risk-neutral quantile functions in the left-tail. What are some other measures that are available at a daily frequency and contain information about the quantile wedge? One candidate is the VIX index, which is defined as

$$VIX_t^2 = \frac{2R_{f,t \rightarrow N}}{N} \left[\int_0^{F_t} \frac{1}{K^2} \text{Put}_t(K) dK + \int_{F_t}^{\infty} \frac{1}{K^2} \text{Call}_t(K) dK \right],$$

where F_t is the forward price on the S&P500, and $\text{Put}_t(K)$ (resp. $\text{Call}_t(K)$) is the put (resp. call) option price on the S&500 with strike K . Martin (2017) shows that

Table F2: **Summary statistics of lower bound**

Horizon	τ	Mean	Median	Std. dev.	Min	Max
30 days	0.05	0.92	0.63	1.07	0.08	24.38
	0.1	0.70	0.45	0.87	0.06	12.22
	0.2	0.47	0.25	0.74	0.04	10.93
60 days	0.05	1.81	1.31	1.67	0.10	19.23
	0.1	1.71	1.19	1.66	0.25	19.89
	0.2	1.14	0.69	1.50	0.12	23.57
90 days	0.05	2.65	2.02	2.02	0.02	18.63
	0.1	2.86	2.12	2.32	0.04	24.47
	0.2	1.97	1.22	2.33	0.26	28.92

Note: This table reports summary statistics of the lower bound, $LB_{t,\tau} = CLB_{t,\tau}/\tilde{f}_t(\tilde{Q}_{t,\tau})$, in (6.13) at different time horizons and different quantile levels over the sample period 2003-2021. All statistics are in percentage point.

VIX measures risk-neutral entropy

$$VIX_t^2 = \frac{2}{N} \tilde{L}_t \left(\frac{R_{m,t \rightarrow N}}{R_{f,t \rightarrow N}} \right),$$

where entropy is defined as $\tilde{L}_t(X) := \log \tilde{\mathbb{E}}_t[X] - \tilde{\mathbb{E}}_t[\log X]$. Entropy, just like variance, is a measure of spread in the distribution. However, entropy places more weight on left-tail events than variance, since entropy places more weight on out-of-the money put options. As such, VIX is a natural candidate to explain potential differences between $Q_{t,\tau}$ and $\tilde{Q}_{t,\tau}$. Second, the Chicago Board Options Exchange provides daily data on VIX at the 30-day horizon.

Table F3 shows the result of the quantile regression

$$Q_{t,\tau}(R_{m,t \rightarrow N}) - \tilde{Q}_{t,\tau}(R_{m,t \rightarrow N}) = \beta_0(\tau) + \beta_1(\tau)LB_{t,\tau} + \beta_{VIX}(\tau)VIX_t. \quad (F.1)$$

We see that β_{VIX} is marginally significant in the left-tail. In contrast, $\beta_1(\tau)$ is even more significant compared to Table 5. Furthermore, the explanatory power of the model that only includes VIX is lower compared to the model that only includes $LB_{t,\tau}$ (see Table 5).

As a second robustness check, I consider how well the direct quantile forecast, $\hat{Q}_{t,\tau} = \tilde{Q}_{t,\tau} + LB_{t,\tau}$, compares to the VIX forecast. Since $\hat{Q}_{t,\tau}$ does not require any parameter estimation, this exercise is a measure of out-of-sample performance. However, VIX does not directly measure $Q_{t,\tau}$ and hence I use an expanding window to

Table F3: **Quantile regression using Lower Bound and VIX**

	$\hat{\beta}_0(\tau)$	$\hat{\beta}_1(\tau)$	$\hat{\beta}_{\text{VIX}}(\tau)$	$R^1(\tau)[\%]$	$R^1(\tau)[\%]$ (VIX only)
$\tau = 0.05$	-0.20 (1.889)	10.09 (0.319)	-0.30 (0.130)	6.34	5.51
$\tau = 0.1$	-0.35 (1.313)	5.06 (0.302)	-0.22 (0.089)	3.41	2.84
$\tau = 0.2$	-0.28 (0.955)	3.62 (0.256)	-0.25 (0.068)	0.61	0.18

Note: This table reports the QR estimates of (F.1) over the 30-day horizon. The sample period is 2003-2021, standard errors are shown in parentheses and calculated using SETBB with a block length equal to the forecast horizon. $R^1(\tau)$ denotes the goodness-of-fit measure (2.5). The last column denotes the goodness-of-fit in the model that only uses VIX as covariate. The standard error and point estimate of β_0 is multiplied by 100 for readability.

obtain the VIX benchmark: $\hat{Q}_{t,\tau}^{\text{VIX}} := \hat{\beta}_0(\tau) + \hat{\beta}_1(\tau)\text{VIX}_t$. Finally, I use the following out-of-sample metric to compare both forecasts

$$R_{\text{oos}}^1(\tau) = 1 - \sum_{t=500}^T \rho_{\tau}(R_{m,t \rightarrow N} - \hat{Q}_{t,\tau}) / \sum_{t=500}^T \rho_{\tau}(R_{m,t \rightarrow N} - \hat{Q}_{t,\tau}^{\text{VIX}}).$$

Notice that $R_{\text{oos}}^1(\tau) > 0$, if $\hat{Q}_{t,\tau}$ attains a lower error than $\hat{Q}_{t,\tau}^{\text{VIX}}$. This exercise is more ambitious, since $\hat{Q}_{t,\tau}^{\text{VIX}}$ makes use of in-sample information. Nonetheless, Figure F5 shows that $\hat{Q}_{t,\tau}$ outperforms the VIX predictor at all percentiles.

Figure F6 performs a similar exercise in the right-tail, but using $\tilde{Q}_{t,\tau}$ instead of $\hat{Q}_{t,\tau}$, since Table 2 shows that the risk-neutral quantile is a good approximation to $Q_{t,\tau}$ in the right-tail. We see that $\tilde{Q}_{t,\tau}$ outperforms $\hat{Q}_{t,\tau}^{\text{VIX}}$ at all quantile levels. Hence, the risk-neutral approximation in the right-tail is more accurate than using the in-sample VIX measure.

F.3 Measurement Error Bias in Quantile Regression

In the empirical application, we have to estimate $\tilde{Q}_{t,\tau}$, $\tilde{f}(\cdot)$ and $\text{CLB}_{t,\tau}$ from market data. Therefore, the estimated coefficients in the quantile regression are biased due to measurement error in the covariate. I present simulation evidence which shows that the bias is small in finite samples.

The setup is as follows. I simulate returns and option prices according to a




Figure F5: **Out-of-sample forecast using risk-adjusted quantile with VIX benchmark.** This figure shows the cumulative out-of sample $R^1(\tau)$, defined as $R^1_{oos}(\tau) = 1 - \sum_{t=500}^T \rho_\tau(R_{m,t \rightarrow N} - \hat{Q}_{t,\tau}) / \sum_{t=500}^T \rho_\tau(R_{m,t \rightarrow N} - \hat{Q}_{t,\tau}^{VIX})$, where $\hat{Q}_{t,\tau} = \hat{Q}_{t,\tau} + \text{LB}_{t,\tau}$, $\hat{Q}_{t,\tau}^{VIX} = \hat{\beta}_0(\tau) + \hat{\beta}_1(\tau) \cdot \text{VIX}_t$, and $\hat{\beta}_0(\tau), \hat{\beta}_1(\tau)$ are the regression estimates from a quantile regression of $R_{m,t \rightarrow N}$ on VIX_t , using data only up to time t . The horizon is 30 days and the QR estimates are dynamically updated using an expanding window over the period 2003–2021. The initial sample uses 500 observations.

discretized version of the [Black and Scholes \(1973\)](#) model:

$$R_{m,t \rightarrow N} = \exp\left(\left(\mu_t - \frac{1}{2}\sigma_t^2\right)N + \sigma_t\sqrt{N}Z_{t+N}\right), \quad Z_{t+N} \sim \mathcal{N}(0, 1) \quad (\text{F.2})$$

$$\sigma_t \sim \text{Unif}[0.05, 0.35]$$

$$\mu_t \sim \text{Unif}[-0.02, 0.2].$$

The return distribution under risk-neutral dynamics is given by

$$\tilde{R}_{m,t \rightarrow N} = \exp\left(\left(r_t - \frac{1}{2}\sigma_t^2\right)N + \sigma_t\sqrt{N}Z_{t+N}\right) \quad (\text{F.3})$$

$$r_t \sim \text{Unif}[0, 0.03]. \quad (\text{F.4})$$

I calculate the lower bound assuming a return horizon of 90 days. As in the

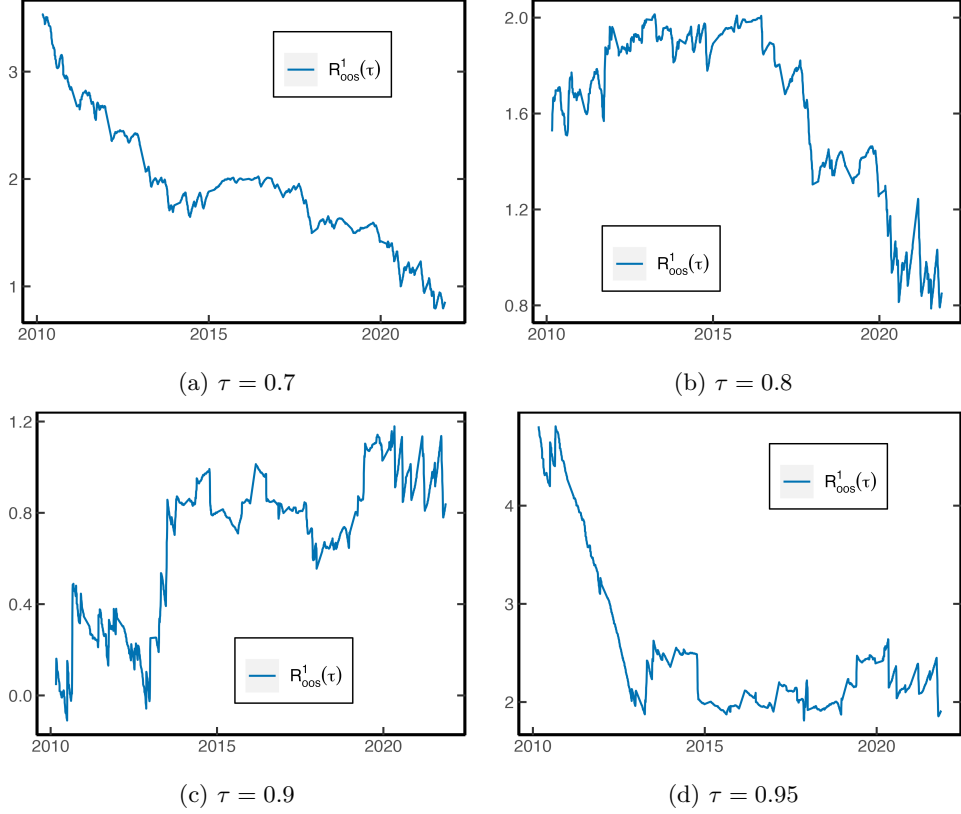


Figure F6: **Out-of-sample forecast using risk-neutral quantile with VIX benchmark.** This figure shows the cumulative out-of sample $R^1(\tau)$, defined as $R^1_{oos}(\tau) = 1 - \sum_{t=500}^T \rho_\tau(R_{m,t \rightarrow N} - \tilde{Q}_{t,\tau}) / \sum_{t=500}^T \rho_\tau(R_{m,t \rightarrow N} - \hat{Q}_{t,\tau}^{VIX})$, where $\hat{Q}_{t,\tau}^{VIX} = \hat{\beta}_0(\tau) + \hat{\beta}_1(\tau) \cdot VIX_t$, and $\hat{\beta}_0(\tau), \hat{\beta}_1(\tau)$ are the regression estimates from a quantile regression of $R_{m,t \rightarrow N}$ on VIX_t , using data only up to time t . The horizon is 30 days and the QR estimates are dynamically updated using an expanding window over the period 2003–2021. The initial sample uses 500 observations.

empirical application, I assume that options with an exact 90-day maturity are not available, but instead we observe options with maturity 85 and 97 days. I generate a total of 1,000 options every time period with maturities randomly sampled from 85 and 97 days.³⁹ These numbers are roughly consistent with the latter part of the empirical sample. The procedure is repeated for a total of 1,000 time periods. For the entire sample, I compare the estimated and analytical lower bound, which are given by respectively

$$LB_{t,\tau}^e := \hat{\tilde{Q}}_{t,\tau} + \frac{\widehat{CLB_{t,\tau}}}{\tilde{f}_t(\hat{\tilde{Q}}_{t,\tau})}$$

$$LB_{t,\tau}^a := \tilde{Q}_{t,\tau} + \frac{CLB_{t,\tau}}{\tilde{f}_t(\tilde{Q}_{t,\tau})}.$$

³⁹So on average there will 500 options with maturity 85 days and 500 with maturity 97 days.

The hats signify that the risk-neutral quantile, PDF and CDF lower bound are estimated from the available options at time t , using the procedure in Appendix C.2. The terms in $LB_{t,\tau}^a$ are obtained from the known analytical expression of the risk-neutral distribution (recall (F.3)). I then use QR to estimate the models

$$Q_{t,\tau} = \hat{\beta}_0(\tau) + \hat{\beta}_{1,e}(\tau)LB_{t,\tau}^e$$

$$Q_{t,\tau} = \hat{\beta}_0(\tau) + \hat{\beta}_{1,a}(\tau)LB_{t,\tau}^a.$$

I use the ratio $\hat{\beta}_{1,e}/\hat{\beta}_{1,a}$ to measure the relative bias in the sample. This experiment is repeated 500 times to get a distribution of the relative bias. Figure F7 shows boxplots of the bias for several quantiles. We see that the relative bias is very small and centered around 1. Hence, the error in measurement problem resulting from estimating the lower bound is limited in this case.




Figure F7: **Bias in QR resulting from measurement error.** This boxplot shows the relative bias in the quantile regression estimate as a result of measurement error.

G Additional Figures




Figure G8: **Highest existing risk-neutral moment for 30-day returns.** This figure shows $p_t^* := \sup\{p : \tilde{\mathbb{E}}_t(R_{m,t \rightarrow N}^p) < \infty\}$ over time, where $R_{m,t \rightarrow N}$ represents the 30-day return. p_t^* is calculated from the moment formula of Lee (2004), $p_t^* = \frac{1}{2\beta_R} + \frac{\beta_R}{8} + \frac{1}{2}$, where $\beta_R = \limsup_{x \rightarrow \infty} \frac{\sigma_{IV}^2(x)}{|x|/N}$, $\sigma_{IV}(x)$ is the implied volatility at log-moneyness $x = \log(K/(e^{rN}S_0))$, and $N = 30/365$ is the time horizon. β_R is estimated from the call option with highest available strike price. The figure is smoothed using a 30-day moving average.