

Quantile Approach to Asset Pricing Models

Tjeerd de Vries *

This Version: December 21, 2021

Abstract

This paper develops a generalization of the Hansen-Jagannathan bound that incorporates information beyond the mean and variance of returns. The resulting bound compares the physical and risk-neutral distribution for every τ -quantile. An empirical application with S&P500 return data shows that the new bound is stronger than the Hansen-Jagannathan bound for small values of τ . This feature is consistent with the rare disaster models, but not with the long run risk model. Using a semiparametric approach, I extend this finding using conditioning information and document that disaster risk is time-varying. I also propose a new measure of quantile forecastability and show that many stylized facts about the equity premium carry over to the quantile setting.

Keywords: Asset pricing, Stochastic discount factor, Quantile methods

JEL Codes: G13, G17, C14, C22

*Department of Economics, University of California San Diego. Email: tjdevrie@ucsd.edu.
I would like to thank Alexis Toda, Allan Timmermann, James Hamilton, Xinwei Ma, Yixiao Sun, Rossen Valkanov and Brendan Beare for useful feedback. All errors are my own.

1 Introduction

Nonparametric methods are useful for analyzing misspecification of asset pricing models in the macro-finance literature (Hansen and Jagannathan, 1991; Snow, 1991; Stutzer, 1995; Bansal and Lehmann, 1997; Alvarez and Jermann, 2005; Backus et al., 2014; Liu, 2020). The underlying theme of these methods is to estimate a statistic of the observable asset returns, such as the Sharpe ratio, and use this statistic to bound unobservable moments of the stochastic discount factor (SDF). The seminal paper of Hansen and Jagannathan (1991) shows how the Sharpe ratio leads to a lower bound on the volatility of the SDF. Hansen and Jagannathan (1991) (henceforth HJ) conclude that any SDF needs to be sufficiently volatile to explain historically high Sharpe ratios. This finding poses a challenge to consumption based asset pricing models, since historical consumption growth in the US is much less volatile than asset returns. A typical problem is that a high level of risk aversion is required in order to overcome the HJ bound, which contradicts other stylized facts of the data.

However, since the HJ bound uses only the mean and variance of excess returns, there might be other information contained in the data that can be exploited to obtain sharper conclusions about the SDF volatility and misspecification of asset pricing models. This paper fills that gap by developing moment bounds on the SDF which use information from the entire distribution of asset returns, not just mean and variance. The new bound on the SDF volatility is motivated as a Sharpe ratio on digital put options, that is, a derivative contract paying out one dollar if the stock falls below a specified strike price. By varying the strike, we obtain a continuum of bounds that compare the physical and risk-neutral distribution for every τ -quantile, where $\tau \in (0, 1)$. For this reason, the new bound is referred to as the quantile bound.

The benefits of this approach are threefold. (i) The quantile bound is valid and performs well even if the distribution of returns is heavy tailed (in contrast to the HJ bound); a point I illustrate in Section 2.3. (ii) Different consumption based asset pricing models imply different shapes of the quantile bound. In particular, some models, like the long-run risk model of Bansal and Yaron (2004), predict that the HJ bound is stronger than the quantile bound. I find counter evidence to this implication in the data (Section 4). I go on to argue that this

finding can be interpreted as model misspecification. Specifically, failure to incorporate disaster risk implies counterfactually high levels of risk-aversion in order to render the quantile bound stronger than the HJ bound. (iii) Quantiles of the left tail of a distribution provide a cleaner interpretation of disaster risk than the mean or variance. I show that introducing disaster risk to an asset pricing model not only increases the market Sharpe ratio, but it also induces a peak in the quantile bound for small τ . I find support for this hypothesis in the data. This finding lends further support to the idea that a fear of disasters is driving expected returns, in contrast to, for example, changing expectations about economic growth as predicted by the long-run risk model.

The analysis above exploits the difference between the unconditional physical and risk-neutral quantile, which I refer to as the quantile wedge. Complementary to the unconditional case, we can ask how the quantile wedge evolves using conditional information. In contrast to the conditional equity premium, the literature offers little guidance on the predictors of this quantile wedge. In addition, the out-of-sample performance is much less understood and typical notions, such as the out-of-sample R_{os}^2 of [Campbell and Thompson \(2008\)](#), are no longer applicable. A second contribution of this paper is to show how these ideas can be incorporated in the quantile setting. Specifically, I propose a clean substitute for R_{os}^2 , denoted by $R_{\text{os}}^1(\tau)$, which is a measure of out-of-sample predictability tailored to quantile regression. Moreover, I derive a forward looking regressor based on option data to approximate the latent conditional quantile function. This regressor consists of the sum of two components, namely the risk-neutral quantile and a risk adjustment term, both of which can be calculated using forward looking information at time t . The results in this paper confirm that the risk adjusted regressor performs well in- and out-of-sample.

This finding leads me to interpret the risk adjusted regressor as a high frequency measure of the latent conditional quantile function. I show that there is significant time fluctuation in the left tail of the quantile function, which can be taken as model-free evidence of time varying disaster risk. Additional insight can be gained from this exercise. For example, I observe that both the risk-neutral and physical quantile drop markedly during the great recession. Interestingly, the decrease in the risk-neutral quantile is far larger. This observation allows us

to gauge the influence of two phenomena that are normally indistinguishable: the increase in perceived disaster risk or the increase in risk premia for insurance against a market decline.¹ My results show that an increase in risk premia were the dominant force during the great recession, even though perceived disaster risk also increased at times of financial distress.

I end the paper with another empirical strategy to test conditional implications of asset pricing models regarding the left and right tail of the distribution of returns. Specifically, I argue that the class of disaster risk models implies that the difference between the physical and risk-neutral distribution is more pronounced in the left tail than the right tail. In contrast, the class of long-run risk models posits that the difference is most pronounced around the median. I propose an empirical strategy to test this hypothesis and find that, consistent with the disaster literature, the difference is most pronounced in the left tail. This finding provides further support for the disaster risk model based on conditional information.

1.1 Literature review

[Hansen and Jagannathan \(1991\)](#) derive a nonparametric bound on the SDF volatility and use it to establish a duality relation with the maximum Sharpe ratio. Many researchers followed up with higher order bounds ([Snow, 1991](#); [Almeida and Garcia, 2012](#); [Liu, 2020](#)) and entropy bounds ([Stutzer, 1995](#); [Bansal and Lehmann, 1997](#); [Alvarez and Jermann, 2005](#); [Backus et al., 2014](#)).

This paper advocates the use of quantiles to obtain sharper SDF bounds. Quantiles from the risk-neutral distribution can be inferred from the market through put and call options, which follows from the work of [Ross \(1976\)](#); [Breen and Litzenberger \(1978\)](#). These authors show that options can complete the market and the risk-neutral PDF can be obtained as the second derivative of the call option price curve. This idea was put into practice by [Ait-Sahalia and Lo \(1998, 2000\)](#); [Rosenberg and Engle \(2002\)](#); [Jackwerth \(2000\)](#) to estimate the SDF or risk-aversion nonparametrically. Even though the approach of [Ait-Sahalia and Lo \(1998, 2000\)](#) renders an estimate of the entire SDF, it requires estimating a ratio of two density functions that are hard to estimate, partic-

¹They are indistinguishable since, normally, we cannot estimate conditional disaster risk.

ularly in the tails. The results proved in this paper yield a bound on certain moments of the SDF, but we only require an estimate of a CDF and quantile function, which, in general, can be estimated at a faster rate.

The results in this paper affirm that the highest Sharpe ratio is not attained by the market portfolio. Instead, I find that it is more profitable to engage in (digital) put selling strategies, which is consistent with [Coval and Shumway \(2001\)](#) and [Bates \(2008\)](#). [Bates \(2008\)](#) proposes a structural model to rationalize this finding using crash-averse preferences. I provide an alternative statistical rationale, leveraging on the empirical observation that asset returns are heavy tailed ([Danielsson and De Vries, 2000](#)). In particular, imposing a Pareto distribution on asset returns with a sufficiently fat tail implies the existence of put selling strategies that yield higher Sharpe ratios than a direct investment in the market portfolio. Importantly, this feature is shared by the class of disaster risk models ([Rietz, 1988](#); [Barro, 2006](#); [Wachter, 2013](#)), but not by the class of long-run risk models ([Bansal and Yaron, 2004](#); [Bansal et al., 2012](#)). I thus provide a new way to compare both models based on distributional characteristics beyond the mean and variance.

This paper also connects to the burgeoning literature on using options to obtain forward looking estimates of the equity premium ([Martin, 2017](#); [Chabi-Yo and Loudis, 2020](#)). However, instead of focusing on the conditional expectation of excess returns, I use option data to predict conditional return quantiles. The evaluation and performance of conditional return predictors is well understood in the literature, especially after fundamental contributions of [Goyal and Welch \(2008\)](#) and [Campbell and Thompson \(2008\)](#). I draw on earlier work of [Koenker and Machado \(1999\)](#) to extend the evaluation toolkit to the quantile setting.

The relation between option data and expected market crashes has a long history and features in the work of [Bates \(1991, 2000, 2008\)](#); [Bollerslev and Todorov \(2011\)](#); [Backus et al. \(2011\)](#); [Ross \(2015\)](#), as well as many others. I build on [Chabi-Yo and Loudis \(2020\)](#) to obtain a semiparametric estimator of the conditional quantile function that does not require parameter estimation or calibration. The time variation in this estimator for low quantiles is consistent with the time varying disaster risk models of [Gabaix \(2012\)](#) and [Wachter \(2013\)](#).

The rest of this paper is organized as follows. [Section 2](#) discusses the quantile

bound, Section 3 outlines how to estimate the new bound with actual data and Section 4 describes the main empirical results. Section 5 extends the empirical results using conditioning information. Finally, Section 6 concludes.

2 Quantile bound on the SDF volatility and model implications

2.1 Notation

Let R be the return on any tradable asset and their portfolios. The assumption of no arbitrage and the law of one price imply the existence of a (positive) stochastic discount factor (SDF), denoted by M , such that the following relation holds

$$\mathbb{E}[MR] = 1. \quad (2.1)$$

The expectation in (2.1) is taken with respect to a probability measure \mathbb{P} , frequently referred to as the physical measure, which represents the true probability distribution of the data generating process. Alternatively, the relation in (2.1) can be reformulated in terms of the risk-neutral measure. That is, there exists a probability measure $\tilde{\mathbb{P}}$, equivalent to \mathbb{P} , such that²

$$\tilde{\mathbb{E}}[R] = 1/\mathbb{E}[M]. \quad (2.2)$$

The expectation on the left in (2.2) is taken with respect to $\tilde{\mathbb{P}}$. Throughout the paper, I use tilde (\sim) to denote quantities that are calculated under the risk-neutral measure $\tilde{\mathbb{P}}$. If we assume the existence of a risk-free asset with return R_f , then $\mathbb{E}[M] = 1/R_f$ and we uncover the familiar fact in (2.2) that the risk-neutral measure absorbs risk premia. Mathematically, the connection between (2.1) and (2.2) follows, since $M/\mathbb{E}[M]$ is the Radon-Nikodym derivative of the measures \mathbb{P} and $\tilde{\mathbb{P}}$. The SDF can potentially depend on many state variables. To avoid having to specify or estimate these state variables, I work with the projected SDF

$$M = \mathbb{E}[\mathfrak{M}|R],$$

²Two probability measures $\mathbb{P}, \tilde{\mathbb{P}}$ are said to be equivalent whenever $\mathbb{P}(A) = 0 \iff \tilde{\mathbb{P}}(A) = 0$.

where \mathfrak{M} is the SDF that depends on all the state variables. The projected SDF has the same pricing implications for contingent claims written on return R for which we have data (Cochrane, 2005, pp. 66–67).³

2.2 Quantile bound

I now prove a new bound on the SDF volatility that elucidates how any difference between the physical and risk-neutral distribution leads to a volatile SDF. Before stating the bound, I introduce $\tilde{F}(\cdot)$ (resp. $F(\cdot)$) to denote the risk-neutral (resp. physical) CDF and write $\tilde{Q}_\tau(R)$ to denote the risk-neutral τ -quantile of the return R . I omit (R) and write \tilde{Q}_τ , whenever the dependence on R is clear from the context. By definition, the risk-neutral quantile function satisfies

$$\tilde{F}(\tilde{Q}_\tau(R)) := \tilde{\mathbb{P}}(R \leq \tilde{Q}_\tau) = \tau.$$

For ease of notation, I also define $\phi(\tau) := F(\tilde{Q}_\tau(R))$, which can be interpreted as the ordinal dominance curve of the measures \mathbb{P} and $\tilde{\mathbb{P}}$ (Hsieh et al., 1996). The new SDF volatility bound can now be stated as follows.

Theorem 2.1 (Quantile bound). *For any $\tau \in (0, 1)$, we have*

$$\frac{\sigma(M)}{\mathbb{E}[M]} \geq \frac{|\tau - \phi(\tau)|}{\sqrt{\phi(\tau) \times (1 - \phi(\tau))}}. \quad (2.3)$$

If a risk-free asset exists, then $\mathbb{E}[M] = 1/R_f$ and (2.3) simplifies to

$$\sigma(M) \geq \frac{|\tau - \phi(\tau)|}{\sqrt{\phi(\tau) \times (1 - \phi(\tau))} R_f}. \quad (2.4)$$

Proof. See Appendix A.1. ■

I refer to Theorem 2.1 as the quantile bound, since a key ingredient in (2.3) is the risk-neutral quantile function. If $\mathbb{P} = \tilde{\mathbb{P}}$, agents are risk-neutral and the dominance curve evaluates to $\phi(\tau) = \tau$. In that case the quantile bound degenerates to zero.

Theorem 2.1 makes precise the sense in which any discrepancy between the physical and risk-neutral distribution induces SDF volatility. Compare this to

³Formally, M is a measurable function of R , but I avoid denoting this dependence explicitly to simplify notation.

the classical HJ bound:

$$\sigma(M) \geq \frac{|\mathbb{E}(R) - R_f|}{\sigma(R)R_f}. \quad (2.5)$$

The lower bound in (2.5) tells us that any excess return induces SDF volatility. Essentially, (2.5) 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. One can imagine that other statistics of the data can be exploited as well to learn something about the SDF. This is precisely what (2.3) does, by comparing the physical and risk-neutral distribution for every τ -quantile. If there is no priced jump or stochastic volatility risk, $\phi(\tau)$ is determined by the equity premium and no information is gained beyond (2.5). On the other hand, if there are risk premia for jumps or stochastic volatility, the measures \mathbb{P} and $\tilde{\mathbb{P}}$ differ both in shape and location (Broadie et al., 2009), which implies that $\phi(\tau)$ contains information that is not captured by (2.5).

Moreover, the quantile bound is well defined regardless of any moment restrictions on the returns and thus robust to fat-tails. Both fat-tails and skewness are essential features of financial return data (Martin and Gao, 2021) and hence the quantile bound could offer useful information beyond that contained in the HJ bound, which is confined to the mean-variance paradigm.

The Sharpe ratio on the right hand side of (2.5) summarizes the risk-return trade-off of a mean-variance optimizing agent. Analogously, the quantile bound can be understood as the Sharpe ratio of insurance against crash risk for small quantiles. Consider a security that pays out \$1, whenever the asset return is below $\tilde{Q}_\tau(R)$, for some small τ . The price of such a security is given by

$$\mathbb{E}[M] \tilde{\mathbb{E}} \left[\mathbb{1} \left(R \leq \tilde{Q}_\tau(R) \right) \right] = \mathbb{E}[M] \tau. \quad (2.6)$$

Similarly, the (discounted) expected return is

$$\mathbb{E}[M] \mathbb{E} \left[\mathbb{1} \left(\tilde{Q}_\tau(R) \right) \right] = \mathbb{E}[M] \phi(\tau). \quad (2.7)$$

And the risk associated to this investment is given by

$$\sigma \left(\mathbb{1} \left(R \leq \tilde{Q}_\tau(R) \right) \right) = \sqrt{\phi(\tau)(1 - \phi(\tau))}. \quad (2.8)$$

Combining (2.6), (2.7) and (2.8) to form the Sharpe ratio we recover (2.3).

The interpretation above sheds light on the economic interpretation of the two bounds. For example, the HJ bound is used to motivate the equity premium puzzle. That is, standard consumption based asset pricing models have difficulty to overcome the HJ bound for reasonable levels of risk aversion, due to historical high returns on the market portfolio (Cochrane, 2005). The disaster risk model of Rietz (1988) and Barro (2006) argues that high Sharpe ratios result from investors' fear of extreme (negative) shocks to consumption. Following this reasoning, we would expect high Sharpe ratios for small quantiles in (2.3) due to the crash risk interpretation sketched above. In this sense, (2.3) is a more direct measure of the influence of disasters on the behavior of the SDF. I analyze this further in Section 2.4.

2.3 An illustrative example

To illustrate how the quantile bound in Theorem 2.1 compares to the HJ bound, I consider a setup with heavy-tailed returns. For ease of exposition, I assume that a risk-free asset exists with return R_f . Let $U \sim \mathbf{UNIF}[0, 1]$ (Uniform distribution on $[0, 1]$) and consider the following specification

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

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

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

It follows from the setup that returns are Pareto distributed.

Proposition 2.2.

- (i) Under physical measure \mathbb{P} , the distribution of returns is Pareto: $R \sim \mathbf{PAR}(B, 1/\beta)$.

(ii) The Sharpe ratio on the asset return is given by

$$\frac{\mathbb{E}[R] - R_f}{\sigma(R)} = \frac{\frac{B}{1-\beta} - \frac{\alpha+1}{A}}{\sqrt{\frac{B^2}{1-2\beta} - \left(\frac{B}{1-\beta}\right)^2}}. \quad (2.9)$$

Proof. See Appendix A.2. ■

I show that the quantile bound in this environment can be stronger than the HJ bound under two different calibrations. To understand the intuition behind this result, I provide an explicit expression for the quantile bound that shows the quantile bound is independent of the Pareto tail parameter β .

Proposition 2.3. *In the setup described above, the following properties hold:*

(i) Under $\tilde{\mathbb{P}}$, $R \sim \mathbf{PAR}\left(B, \frac{\alpha+1}{\beta}\right)$.

(ii) The quantile bound depends only on the (left) tail index α of M . In particular,

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

(iii) If $\beta \uparrow \frac{1}{2}$, then the HJ bound converges to 0.

Proof. See Appendix A.2. ■

Properties (ii) and (iii) provide some intuition when the quantile bound is stronger than the HJ bound. Namely, heavier tails of the distribution of R (as measured by β) lead to a lower Sharpe ratio. However, the quantile 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 quantile bound may fare better. Moreover, we do not need to impose any restrictions on the parameter space to calculate the quantile bound, whereas the HJ bound requires $\beta < 1/2$. However, the 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).

I now calibrate the model in two different ways to illustrate the difference between the quantile and HJ bound. The first calibration is targeted to match

some of the salient features of the US market return. To match the typical Pareto exponent for the US market return, I set $\beta = 1/3$. For simplicity, assume $R_f = 1$ and to match the observed equity premium of 8%, I pick $B = 1.08 \times (1 - \beta) = 0.72$. The implied return volatility is far above typical estimates (16% in [Cochrane \(2005\)](#)), but since this example is provided to gain intuition the discrepancy is ignored. Equations (A.5) and (A.6) are used to solve for the parameters A, α from the SDF distribution. In the alternative calibration, I set $\beta = 1/2.2$ (heavier tails, but still finite variance) and pick B again to match the equity premium. Once more, (A.5) and (A.6) are used to solve for A, α . Table 1 summarizes the resulting parameter values for reference, together with the corresponding Sharpe ratio and SDF volatility.

Table 1: Model calibration

	A	α	B	β	$\sigma(R)$	Sharpe ratio	$\sigma(M)$
Calibration 1	1.19	0.19	0.72	0.33	0.62	0.13	0.16
Calibration 2	1.11	0.11	0.59	0.45	1.63	0.05	0.10

Note: Calibration of SDF model with Pareto returns. Both calibrations impose an equity premium of 8% and (gross) risk-free rate $R_f = 1$. $\sigma(R)$ denotes the return volatility and $\sigma(M)$ the SDF volatility.

Figure 1 shows the HJ and quantile bound for the two different calibrations. One can see that the supremum of the quantile bound exceeds the HJ bound in both cases. The display on the right shows that the performance of the quantile bound is better whenever the distribution of asset returns is more heavy tailed. In the left display of Figure 1, the supremum of the quantile bound is slightly stronger than the HJ bound, but only marginally so. In contrast, the display on the right of Figure 1 highlights that there is a range of values for which the quantile bound is stronger than the HJ bound, owing to the fatter tails of the return distribution.

2.4 Quantile and HJ bound for common asset pricing models

In this section, I compare the tightness of the quantile bound in Theorem 2.1 to the HJ bound using common asset pricing models. The tightness is of interest, since some asset pricing models imply that the HJ bound is always tighter

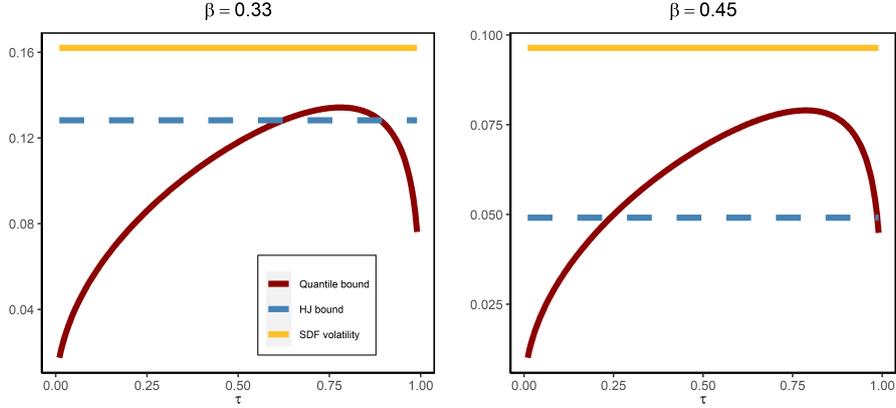


Figure 1: Plots of the quantile bound (in red), HJ bound (blue), true SDF volatility (yellow) and estimated quantile bound (yellow) for different values of β .

than the quantile bound. For other models, this feature is true under common parameter calibration. Since real data in Section 4 show that the quantile bound is significantly stronger than the HJ bound, this finding can be taken as evidence against such models. Appendix D contains similar results in this direction using other well known asset pricing bounds.

Example 2.1 (CAPM). The Capital Asset Pricing Model (CAPM) specifies the SDF as

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

Here, R_m denotes the return on the market portfolio. Since the HJ bound is derived by applying the Cauchy-Schwarz inequality to $\text{COV}(R_m, M)$, the inequality binds if M is a linear combination of R_m . Hence, under CAPM, the HJ bound is always stronger than the quantile bound regardless of the distribution of R_m .

For the following two examples I need Stein's Lemma (Cochrane, 2005, p. 163):

Lemma 2.4 (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*

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

Example 2.2 (Joint normality). Suppose that M and R are jointly normally distributed. This 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 2.3). The proof of the quantile bound in Theorem 2.1 gives the following identity

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

By an approximation argument, Stein's lemma still applies with $g(x) = \mathbb{1}(x \leq k)$ and $g'(x) = \delta_k(x)$ (Dirac delta function). Therefore,

$$\left| \text{COV} \left(\mathbb{1} \left(R \leq \tilde{Q}_\tau \right), M \right) \right| = f(\tilde{Q}_\tau) |\text{COV}(R, M)|. \quad (2.10)$$

Here, $f(\cdot)$ is the marginal density of R . Standard SDF properties also yield the well known identity

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

To evaluate the strengths of the quantile and HJ bound, consider the relative efficiency

$$\begin{aligned} \frac{\text{HJ bound}}{\text{Quantile bound}} &= \frac{\frac{|\mathbb{E}[R] - R_f|}{\sigma(R)R_f}}{\frac{|\tau - \phi(\tau)|}{\sqrt{\phi(\tau)(1 - \phi(\tau))}R_f}} \\ &\stackrel{(2.10)}{=} \frac{\sqrt{\phi(\tau)(1 - \phi(\tau))}}{\sigma(R)f(\tilde{Q}_\tau)}. \end{aligned} \quad (2.11)$$

To see that the HJ bound is always stronger than the quantile bound, minimize (2.11) with respect to τ . Temporarily write $x = \tilde{Q}_\tau$ and consider

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

Minimizing $\Gamma(x)$ is equivalent to minimizing (2.11) and first order conditions imply that the optimal x^* satisfies

$$[f(x^*) - 2F(x^*)f(x^*)]f(x^*)^2 - 2f(x^*)f'(x^*)[F(x^*)(1 - F(x^*))] = 0. \quad (2.12)$$

Since f, F are the respective PDF and CDF of the normal random variable R ,

it follows that $f'(\mu_R) = 0$ and $F(\mu_R) = 1/2$, where μ_R is the mean of R . As a result, (2.12) holds when $\tilde{Q}_{\tau^*} = x^* = \mu_R$. For this choice, $\phi(\tau^*) = \mathbb{P}(R \leq \tilde{Q}_{\tau^*}) = 1/2$ and $f(\tilde{Q}_{\tau^*}) = 1/\sqrt{2\pi\sigma_R^2}$. Therefore, (2.11) obeys the bound

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

Hence, the HJ bound is always better in a model where the SDF and wealth portfolio are assumed to be jointly normal.

Example 2.3 (Joint lognormality). Let Z_R and Z_M be standard normal random variables with correlation ρ and consider the specification

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

where λ governs the time scale. Simple algebra shows that the no-arbitrage condition $\mathbb{E}[RM] = 1$ is satisfied when $\mu_R - r_f = -\rho\sigma_R\sigma_M$. It is hard to find an analytical solution for the relative efficiency between the HJ and quantile 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.3, where I prove that

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

This expression is independent of μ_R . An application of l'Hôpital's rule reveals that the relative efficiency converges to $\sqrt{2\pi}/2$ if $\lambda \rightarrow 0^+$. This is the same relative efficiency in Example 2.2, which is unsurprising as the linearization becomes exact in the limit as $\lambda \rightarrow 0^+$. This ratio is less than 1 if $\sigma \geq 0.91$ and $\lambda = 1$. In practice, annualized market return volatility is about 16%, which means that the HJ bound is stronger than the quantile bound under any reasonable parameterization if the SDF and asset return are lognormal.

Example 2.4 (Disaster risk). The disaster risk model of Rietz (1988) and Barro (2006) posits that risk-premia are driven by extreme events that affect consumption growth. I follow the specification in Backus et al. (2011), who assume that the representative agent has power utility and the log pricing kernel

is given by

$$\log M = \log(\beta) - \gamma\Delta c.$$

Innovations in consumption growth are driven by two independent shocks

$$\Delta c = \varepsilon + \eta. \tag{2.13}$$

Here, $\varepsilon \sim N(\mu, \sigma^2)$ and

$$\eta|(J = j) \sim N(j\theta, j\nu^2), \quad J \sim \mathbf{Poisson}(\kappa).$$

The interpretation of η is that of a jump component (disaster) which induces negative shocks to consumption growth. κ governs the jump intensity for the Poisson distribution. I use the same calibration as [Backus et al. \(2011\)](#). In line with their paper, the market portfolio is considered as a claim on levered consumption, i.e. an asset that pays dividends C^λ . I convert the model implied volatility bounds to monthly units, to facilitate the comparison with the long-run risk model and the empirical bounds obtained in [Section 4](#).

The quantile bound, HJ bound and SDF volatility are illustrated in the left upper panel of [Figure 2](#).⁴ We see that the quantile bound has a sharp peak at $\tau = 0.046$, after which it decreases monotonically. Interestingly, there is a range of τ values for which the quantile bound is sharper than the HJ bound. This is in line with the empirical evidence in [Section 4](#).

The reason for this result can be understood from the upper right panel of [Figure 2](#), which shows the physical and risk-neutral distribution of return on equity. The right tail of the distributions are not shown, since they are virtually indistinguishable in that region. The risk-neutral distribution displays a heavy left tail, owing to the implied disaster risk embedded in the SDF. As a result, 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 disaster is so low that the risk associated to selling such insurance is limited.

The following example uses time t conditional information. I denote the

⁴The bounds can be calculated analytically, see [Appendix C.1](#) for the details.

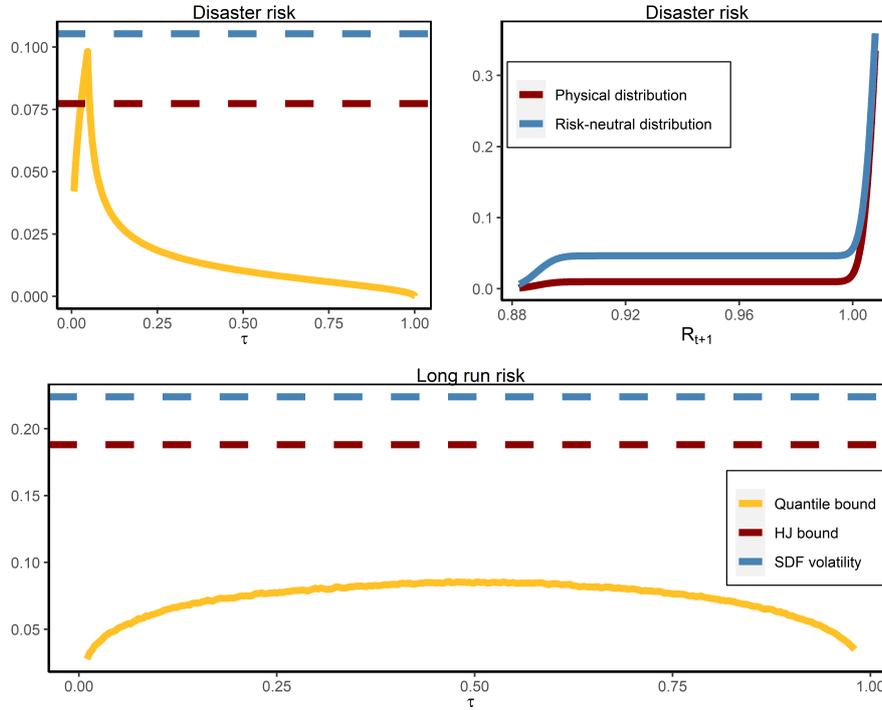


Figure 2: The upper left panel compares the HJ bound and quantile bound in disaster risk model and the right upper panel shows the physical and risk-neutral distribution. The bottom panel shows the HJ and quantile bound for the long run risk model. The bounds are in monthly units.

SDF from time t to $t + 1$, by M_{t+1} . The subscript refers to a random variable that is realized at time $t + 1$, conditioned on time t .

Example 2.5 (Long-run risk). The long-run risk (LRR) model of [Bansal and Yaron \(2004\)](#) posits that consumption growth is driven by a small and persistent component that captures long-run risk. Moreover, the existence of a representative agent with [Epstein and Zin \(1989\)](#) recursive preferences is assumed. After calibration, this model is successful in matching many of the salient features of the US market return data. I consider the extended model of [Bansal et al. \(2012\)](#), which allows for correlation between consumption growth shocks and

dividend growth. In particular, the following dynamics are assumed:

$$\begin{aligned}
x_{t+1} &= \rho x_t + \varphi_e \sigma_t e_{t+1} \\
\sigma_{t+1}^2 &= \bar{\sigma}^2 + \nu(\sigma_t^2 - \bar{\sigma}^2) + \sigma_w w_{t+1} \\
\Delta c_{t+1} &= \mu_c + x_t + \sigma_t \eta_{t+1} \\
\Delta d_{t+1} &= \mu_d + \phi x_t + \pi \sigma_t \eta_{t+1} + \varphi \sigma_t u_{d,t+1}.
\end{aligned}$$

Here, Δc_{t+1} and Δd_{t+1} denote log consumption and dividend growth, while σ_t is conditional volatility of log consumption growth. The parameter ρ governs the persistence of long-term risk. The log SDF dynamics follow from the Euler equation and the [Epstein and Zin \(1989\)](#) preferences

$$\log M_{t+1} = \theta \log \beta - \frac{\theta}{\psi} \Delta c_{t+1} + (\theta - 1) r_{c,t+1},$$

where $r_{c,t+1}$ is the continuous return on the consumption asset. I omit further details on the parameter interpretation and calibration approach, which is extensively discussed in [Bansal et al. \(2012\)](#). To compare the HJ bound to the quantile bound, I use the same calibration of parameters as [Bansal et al. \(2012\)](#). The quantile bound, as well as the HJ bound and SDF volatility are estimated by simulation, using 110,000 months, where the first 10,000 observations are dropped as burn-in sample. The results are summarized in the bottom panel of [Figure 2](#). Unlike the disaster risk model, the quantile bound is always weaker than the HJ bound. Moreover, the quantile bound is almost symmetric around $\tau = 0.5$, at which the maximum is attained. Hence, it is not profitable to sell insurance against disaster risk. This contradicts the empirical estimates from [Section 4](#), which suggests that the quantile bound is stronger than the HJ bound and implies it is most profitable to sell insurance against disaster risk.

3 Estimation of the unconditional quantile bound

In this section I discuss the estimation of the (unconditional) quantile bound. I assume the existence of a risk-free asset, so we can use the quantile bound in [\(2.4\)](#). The bound is comprised of three unknowns: the risk-neutral quantile

function $\tilde{Q}_\tau(R)$, the physical probability measure $\mathbb{P}(R \leq x)$ and the unconditional risk-free rate R_f . I follow [Liu \(2020\)](#) and fix R_f at a pre-specified level. This is because the US risk-free rate can be inferred with high precision and its influence is minor compared to the estimation of the first two functions. I sketch the intuition for the function estimates below; a more detailed description is provided in [Appendix B.2](#).

Since our estimation strategy involves returns sampled over time, we have to incorporate conditioning information. I use R_{t+1} to denote the asset return realized at time $t+1$, conditioned on time t information. I use the same notation for the risk-free rate $R_{f,t+1}$, which is assumed to be known at time t .

To estimate the risk-neutral quantile function, I first estimate $\tilde{\mathbb{P}}(R \leq x)$ and use inversion to get an estimate of $\tilde{Q}_\tau(R)$. To make this idea operational, we rely on the following result of [Breedon and Litzenberger \(1978\)](#)

$$\tilde{\mathbb{P}}_t(R_{t+1} \leq K/S_t) = R_{f,t+1} \frac{\partial}{\partial K} \text{Put}_t(K),$$

where $\text{Put}_t(K)$ is the time t price of a European put option with strike K , expiring at time $t+1$ and S_t is the time t stock price of the underlying asset. Hence, with enough put option prices, we can identify the conditional risk-neutral CDF. To obtain an estimate of the unconditional CDF, we average the conditional CDFs over time

$$\tilde{F}_T(x) := \tilde{\mathbb{P}}_T(R \leq x) := \frac{1}{T} \sum_{t=1}^T \tilde{\mathbb{P}}_t(R_{t+1} \leq x).^5 \quad (3.1)$$

Under suitable restrictions on the distribution of returns, we expect this to converge to the unconditional return distribution. An estimate for the unconditional risk-neutral quantile curve is then obtained from

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

It is a nontrivial exercise to obtain solid estimates via this procedure, 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. A detailed

⁵I follow the empirical process literature and use subscript T to denote a functional estimate based on T observations.

description is given in Appendix B.2, using a modification of the procedure proposed by Figlewski (2008).

Secondly, we need to estimate $\mathbb{P}(R \leq x)$. I use a kernel CDF estimator for this function, which ensures that the quantile bound is a smooth function of τ . Imposing a smooth quantile bound improves the approximation in finite samples and mitigates the influence of outliers that would be more pronounced, if, say, we use the discontinuous empirical CDF estimator. The kernel CDF estimator is given by

$$F_T(x) := \mathbb{P}_T(R_{t+1} \leq x) = \frac{1}{T} \sum_{t=1}^T \Phi\left(\frac{x - R_{t+1}}{h}\right),$$

where Φ is the integral of the Epanechnikov kernel and h is the bandwidth. The optimal bandwidth is determined via cross-validation. I denote the estimated dominance curve by $\phi_T(\tau) := F_T(\tilde{Q}_T(\tau))$. Combining the pieces, I obtain the following estimator for the quantile bound

$$\hat{\theta}(\tau) := \frac{|\tau - \phi_T(\tau)|}{\sqrt{\phi_T(\tau) [1 - \phi_T(\tau)]} R_f} \quad \tau \in [\varepsilon, 1 - \varepsilon], \quad (3.2)$$

for some small $\varepsilon > 0$.

4 Empirical application

This Section presents estimates of the quantile bound using a combination of forward looking option data and historical market returns. Based on calibration of the disaster risk model, we provide statistical evidence that the quantile bound renders a stronger bound on the SDF volatility than the HJ bound.

4.1 Data and estimation of the unconditional risk-neutral quantile curve

I use daily option data on the S&P500 from OptionMetrics covering the period 01-01-1996 until 12-31-2019 to estimate the unconditional risk-neutral quantile curve. Before estimating the risk-neutral quantile curve, I use several data cleaning procedures, which are detailed in Appendix B.1. Subsequently, the

unconditional risk-neutral quantile curve is estimated following the steps in Appendix B.2. Figure 3 shows the estimated unconditional quantile curves for several horizons. It is apparent that longer time horizons carry a higher (risk-neutral) probability of up and downswings. This makes intuitive sense, but is not apparent a priori, since the quantile curves document profit and losses over a fixed horizon and not during any time within the horizon.

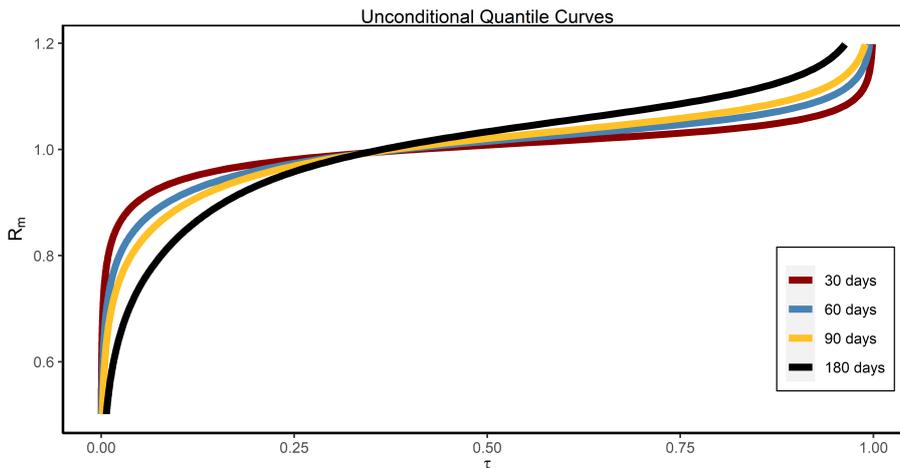


Figure 3: Plots of the unconditional (risk-neutral) quantile function for various maturities.

4.2 Quantile bound for 30-day returns

I now turn to the estimation of the unconditional quantile bound for 30-day returns, using the estimator in (3.2). The unconditional physical CDF estimate is based on *non-overlapping*, historical 30-day returns on the S&P500 index over the period 1996-2019, calculated at the middle of the month. This consists of a total of $T = 288$ return observations. The unconditional risk-neutral quantile function is estimated following the procedure in Appendix B.2, using only the dates at which the historical market returns $R_{m,t+1}$ are calculated; that is, I average over dates t , corresponding to the start of the return period of $R_{m,t+1}$. Finally, I fix the unconditional risk-free rate at $R_f = 1$, which roughly corresponds to the sample average of monthly risk-free rates over the time period 1996-2019.

Figure 4 shows the estimated quantile bound, as well as the HJ bound. Notice that the estimated quantile bound is akin to the quantile bound predicted

by the disaster risk model (Figure 2). This finding corroborates the idea that very high Sharpe ratios can be obtained for selling insurance against crash risk and is consistent with empirical evidence from previous studies that analyze the Sharpe ratios of put selling strategies (see Broadie et al. (2009) and the references therein).

The shape of the quantile bound in Figure 4 is also similar to the direct SDF estimates of Aït-Sahalia and Lo (2000) and Rosenberg and Engle (2002), which are based on the ratio $\zeta_t(R_{m,t+1}) = \tilde{f}_t(R_{m,t+1})/f_t(R_{m,t+1})$. In their context, $1/\zeta_t(R_{m,t+1})$ can be thought of as the expected return from buying an Arrow-Debreu security, which pays out \$1 if the return next period is $R_{m,t+1}$. Aït-Sahalia and Lo (2000, Table 3) report large negative expected returns for small values of $R_{m,t+1}$, implying that selling Arrow-Debreu securities for low returns is profitable. The authors do not further develop the implications for asset pricing models, as I do below.

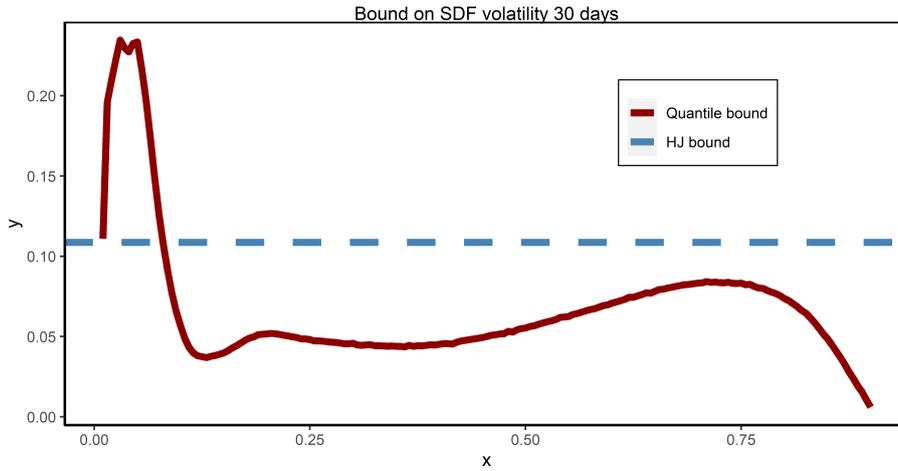


Figure 4: Plot of the quantile bound as function of τ . The solid red line is the estimated quantile bound. The dashed blue line depicts the HJ bound.

4.2.1 Testing the quantile and HJ bound

I now test more formally whether the quantile bound improves upon the HJ bound, using a bootstrap approach. To make this operational, I fix, a priori, the τ -quantile at probability level $\tau = 0.046$, which renders the sharpest bound on the SDF volatility in the disaster risk model (Example 2.4). The resulting

quantile bound is summarized in Table 2, together with the HJ bound. Observe that the quantile bound is more than double that of the HJ bound. To see whether this is a statistically significant difference, I consider the test statistic

$$\mathcal{T}\{R_{m,t+1}\} := \widehat{\theta}(0.046) - \frac{|\bar{R}_m - R_f|}{\widehat{\sigma}(R_m)R_f}. \quad (4.1)$$

The first term on the right hand side denotes the estimated quantile bound (3.2) evaluated at the 0.046-quantile, using the entire time series $\{R_{m,t+1}\}$ of returns. The second term denotes the estimated HJ bound, using \bar{R}_m and $\widehat{\sigma}(R_m)$ as the respective sample mean and standard deviation of $\{R_{m,t+1}\}$. A value of $\mathcal{T}\{R_{m,t+1}\} \geq 0$ indicates that the quantile bound is stronger than the HJ bound.

Since the distribution of (4.1) is hard to characterize, I use block bootstrap to approximate the p -value of the null hypothesis

$$H_0 : \mathcal{T}\{R_{m,t+1}\} \leq 0. \quad (4.2)$$

The block bootstrap is used to generate time indices from which we recreate (with replacement) bootstrapped returns $\{R_{m,t+1}^*\}$. The same bootstrapped time indices are used to estimate the risk-neutral quantile curves and leads to the quantile bound $\widehat{\theta}^*(0.046)$, as well as the HJ bound

$$\frac{|\bar{R}_m^* - R_f|}{\widehat{\sigma}(R_m^*)R_f}.$$

I repeat the bootstrap exercise 100,000 times and calculate the bootstrapped test statistic

$$\mathcal{T}\{R_{m,t+1}^*\} := \widehat{\theta}^*(0.046) - \frac{|\bar{R}_m^* - R_f|}{\widehat{\sigma}(R_m^*)R_f}. \quad (4.3)$$

Finally, the p -value is obtained as the fraction of times $\mathcal{T}(R_{m,t+1}^*) \leq 0$. Table 2 shows that the estimated p -value is 0.0529, suggesting that the quantile bound is stronger than the HJ bound.

4.2.2 Implications

What are the economic implications of this finding? First, in Section 2.4, I argued that under CAPM, the HJ bound is always stronger than the quantile

Table 2: Sample bounds and bootstrap result

R_f	Sample size	HJ bound	τ	Quantile bound	p -value
1	288	0.1087	0.0464	0.2339	0.0529

Note: Test result of (4.2), using 100,000 bootstrap samples. The HJ and quantile bound refer to the sample estimates, using the entire time of observations.

bound. Therefore, the empirical finding that the quantile bound is stronger than the HJ bound is evidence against CAPM. The LRR model of [Bansal et al. \(2012\)](#) can neither reconcile this feature of the data under common parameter calibration. The evidence from [Figure 4](#) and [Table 2](#) supports the view that misspecification of the LRR model stems from the absence of disaster risk. For the LRR model, I establish by simulation that a risk aversion coefficient of 90 is needed for the quantile bound to overcome the HJ bound. Moreover, in the LRR model, the quantile bound is almost symmetric around $\tau = 0.5$ and decreases steeply when τ is close to $\{0, 1\}$, which again contradicts empirical evidence that the highest Sharpe ratios are obtained for selling insurance against disaster risk. The empirical results thus favor the theoretical implications of the disaster risk model.

The observation that put options can be leveraged to yield higher Sharpe ratios than a direct investment in the market portfolio has been noted before in the literature ([Bates, 2008](#); [Broadie et al., 2009](#)). However, what seems to have gone unnoticed is the connection to model misspecification, in the sense that this empirical regularity puts a tight constraint on a model’s risk-aversion coefficient. In particular, this is true for the LRR model or statistical models that impose joint lognormality between the SDF and market return ([Example 2.3](#)). In the most extreme case, for CAPM or models that impose joint normality ([Example 2.2](#)), we cannot even find parameters that can generate the empirical regularity. These results underscore the importance of incorporating information beyond the mean and variance to analyze model misspecification.

5 Conditional relation between physical and risk-neutral quantile

The unconditional quantile bound puts a tight constraint on the SDF volatility due to the apparent disaster risk in the data. This section complements that finding by looking at conditional information. I show how the numerator term in the quantile bound governs the difference between the physical and risk-neutral quantile, up to a first order correction. I interpret this first order correction as a risk adjustment term, which, under some conditions, can be approximated by option data available at time t . Fluctuation in the risk adjustment term for small τ points to time varying disaster risk.

5.1 Lower bound on physical quantile

I start by deriving a market observable lower bound on the difference between the physical and risk-neutral quantile function. Write $F_t(x) := \mathbb{P}_t(R_{t+1} \leq x)$ for the physical distribution function of R_{t+1} conditional on all information available at time t , $f_t(\cdot)$ for the PDF and $Q_{t,\tau}$ for the conditional τ -quantile. As before, a tilde superscript refers to the risk-neutral measure. The idea is to use [von Mises \(1947\)](#) calculus for statistical functionals to find a first order approximation to the physical quantile function. A statistical functional $\varphi : \mathbb{D} \rightarrow \mathbb{E}$ is typically a map between two normed spaces \mathbb{D} and \mathbb{E} . In our application, we take $\varphi(G) = G^{-1}(\tau)$, where G is a CDF and $G^{-1}(\tau)$ is the quantile function. Following [Van der Vaart \(2000, Section 20.1\)](#) and [Serfling \(2009, p. 217\)](#), we have the following identity

$$Q_{t,\tau} - \tilde{Q}_{t,\tau} = \varphi(F_t) - \varphi(\tilde{F}_t) = \varphi'_{\tilde{F}_t}(F_t - \tilde{F}_t) + o\left(\|F_t - \tilde{F}_t\|\right), \quad (5.1)$$

where $\|\cdot\|$ is a norm on a suitable linear space⁶ 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

$$\begin{aligned}\varphi'_{\tilde{F}_t}(F_t - \tilde{F}_t) &:= \lim_{\lambda \downarrow 0} \frac{\varphi \left[(1 - \lambda)\tilde{F}_t + \lambda F_t \right]}{\lambda} \\ &= \frac{\partial}{\partial \lambda} \varphi \left((1 - \lambda)\tilde{F}_t + \lambda F_t \right) \Big|_{\lambda=0}.\end{aligned}\quad (5.2)$$

Heuristically, the Gâteaux derivative in this context can be thought of as measuring the change in the quantile function when we move the distribution of returns from risk-neutral in the direction of the physical distribution. Appendix A.7 shows that, when $\varphi(\cdot)$ is the quantile map, the Gâteaux derivative is given by

$$\varphi'_{\tilde{F}_t}(F_t - \tilde{F}_t) = \frac{\tau - F_t(\tilde{Q}_{t,\tau})}{\tilde{f}_t(\tilde{Q}_{t,\tau})}.\quad (5.3)$$

We now make it our working hypothesis that the remainder term in (5.1) is “small” in the sup-norm, $\|g\|_\infty = \sup_x |g(x)|$.

Working Hypothesis 5.1. *The remainder term in (5.1), $\|F_t - \tilde{F}_t\|_\infty$, can be neglected.*

Remark. The word “small” might be nebulous in this context; all it means is that we assume that the first order approximation in (5.1) is accurate. The assumption that $\|F_t - \tilde{F}_t\|_\infty$ is small is quite natural, since the discussion in Example 2.4 shows that substantial pointwise difference between $F_t(\cdot)$ and $\tilde{F}_t(\cdot)$ leads to near-arbitrage opportunities. Ultimately, whether the Working Hypothesis is reasonable or not is an empirical question. I test this in Section 5.2 and find that the approximation is quite accurate. Additionally, Appendix E shows that Assumption 5.1 finds support in the Black and Scholes (1973) model.

Combining (5.1) and (5.3) in conjunction with Assumption 5.1 renders the approximation

$$Q_{t,\tau} \approx \tilde{Q}_{t,\tau} + \underbrace{\frac{\tau - F_t(\tilde{Q}_{t,\tau})}{\tilde{f}_t(\tilde{Q}_{t,\tau})}}_{\text{risk adjustment}}.\quad (5.4)$$

Observe that the numerator in the risk adjustment term equals the numerator

⁶Formally, 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).

term in the conditional version of the quantile bound. The approximation in (5.4) contains the terms $\tilde{Q}_{t,\tau}$ and $\tilde{f}_t(\tilde{Q}_{t,\tau})$, which are directly observed at time t , as they can be computed using (a variation of) the [Breedeen and Litzenberger \(1978\)](#) formula in (3.1). However, the physical CDF, $F_t(\cdot)$, is unknown and hence (5.4) cannot directly be used to predict $Q_{t,\tau}$.

Under additional assumptions, the numerator term $\tau - F_t(\tilde{Q}_{t,\tau})$ can be bounded with market data. To show this, I link $F_t(\cdot)$ to a risk-neutral covariance term, similar to [Chabi-Yo and Loudis \(2020\)](#). Use the reciprocal of the SDF to pass from physical to risk-neutral measure

$$\begin{aligned} F_t(\tilde{Q}_{t,\tau}) &= \mathbb{E}_t \left[\mathbb{1} \left(R_{t+1} \leq \tilde{Q}_{t,\tau} \right) \right] = \tilde{\mathbb{E}}_t \left[\mathbb{1} \left(R_{t+1} \leq \tilde{Q}_{t,\tau} \right) \frac{\mathbb{E}_t [M_{t+1}]}{M_{t+1}} \right] \\ &= \widetilde{\text{COV}}_t \left[\mathbb{1} \left(R_{t+1} \leq \tilde{Q}_{t,\tau} \right), \frac{\mathbb{E}_t [M_{t+1}]}{M_{t+1}} \right] + \tau. \end{aligned} \quad (5.5)$$

To proceed, assume that $R_{t+1} = R_{m,t+1}$ (the market return). This allows me to get a more explicit expression of the SDF as follows. [Chabi-Yo and Loudis \(2020\)](#) show that in a one-period model with a representative agent who has utility function $u(\cdot)$ and derives utility over final wealth, the SDF is given by

$$\frac{\mathbb{E}_t [M_{t+1}]}{M_{t+1}} = \frac{\frac{u'(W_t x_0)}{u'(W_t x)}}{\tilde{\mathbb{E}}_t \left[\frac{u'(W_t x_0)}{u'(W_t x)} \right]} \quad \text{with } x = R_{m,t+1}, \quad x_0 = R_{f,t+1},$$

and W_t is the agent's initial wealth at time t . Define $f(x) := \frac{u'(W_t x_0)}{u'(W_t x)}$ and use a Taylor expansion around $x = x_0$ to get

$$f(x) = 1 + \sum_{k=1}^{\infty} \theta_k (x - x_0)^k \quad \text{with } \theta_k = \frac{1}{k!} \left(\frac{\partial^k f(x)}{\partial x^k} \right)_{x=x_0}.$$

The θ_k -coefficients depend on the specific utility representation employed, but

are conditionally non random. I substitute the above in (5.5) and obtain

$$\begin{aligned}
& \widetilde{\text{COV}}_t \left[\mathbb{1} \left(R_{m,t+1} \leq \tilde{Q}_{t,\tau} \right), \frac{\mathbb{E}_t [M_{t+1}]}{M_{t+1}} \right] \\
&= \widetilde{\text{COV}}_t \left[\mathbb{1} \left(R_{m,t+1} \leq \tilde{Q}_{t,\tau} \right), \frac{f(R_{m,t+1})}{\tilde{\mathbb{E}}_t(f(R_{m,t+1}))} \right] \\
&= \frac{\sum_{k=1}^{\infty} \theta_k \left(\tilde{\mathbb{E}}_t \left[\mathbb{1} \left(R_{m,t+1} \leq \tilde{Q}_{t,\tau} \right) (R_{m,t+1} - R_{f,t+1})^k \right] - \tau \tilde{\mathbb{E}}_t \left[(R_{m,t+1} - R_{f,t+1})^k \right] \right)}{1 + \sum_{k=1}^{\infty} \theta_k \tilde{\mathbb{E}}_t \left[(R_{m,t+1} - R_{f,t+1})^k \right]}.
\end{aligned} \tag{5.6}$$

In the last line I used $\tilde{\mathbb{E}}_t(\mathbb{1}(R_{m,t+1} \leq \tilde{Q}_{t,\tau})) = \tau$. The upshot is that the (un)truncated moments of the excess market return can be computed from option data (see Appendix A.4). To enhance notation, I follow Chabi-Yo and Loudis (2020) and write

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

Using the more compact notation, we have

$$\widetilde{\text{COV}}_t \left[\mathbb{1} \left(R_{m,t+1} \leq \tilde{Q}_{t,\tau} \right), \frac{\mathbb{E}_t [M_{t+1}]}{M_{t+1}} \right] = \frac{\sum_{k=1}^{\infty} \theta_k \left(\tilde{\mathbb{M}}_{t+1}^{(k)}[\tilde{Q}_{t,\tau}] - \tau \tilde{\mathbb{M}}_{t+1}^{(k)} \right)}{1 + \sum_{k=1}^{\infty} \theta_k \tilde{\mathbb{M}}_{t+1}^{(k)}}. \tag{5.7}$$

Combining Equation (5.7) and (5.5) in (5.4) leads to the first order approximation

$$Q_{t,\tau} \approx \tilde{Q}_{t,\tau} + \frac{1}{\tilde{f}_t(\tilde{Q}_{t,\tau})} \left(\frac{\sum_{k=1}^{\infty} \theta_k \left(\tau \tilde{\mathbb{M}}_{t+1}^{(k)} - \tilde{\mathbb{M}}_{t+1}^{(k)}[\tilde{Q}_{t,\tau}] \right)}{1 + \sum_{k=1}^{\infty} \theta_k \tilde{\mathbb{M}}_{t+1}^{(k)}} \right). \tag{5.8}$$

The right hand side of (5.8) depends on known quantities that can be calculated at time t with option data, except for the unknown parameters θ_k . However, Chabi-Yo and Loudis (2020) show that we can make assumption about θ_k that lead to a lower bound. I adopt the following assumptions from their paper:

Assumption 5.2. $\tilde{M}_{t+1}^{(k)} \leq 0$ if k is odd. Furthermore, for $k_0 \leq R_{f,t+1}$

$$\begin{aligned}\tilde{M}_{t+1}^{(1)}[k_0] &\leq 0, & \tilde{M}_{t+1}^{(2)}[k_0] &\geq 0 \\ \tilde{M}_{t+1}^{(3)}[k_0] &\leq 0, & \tilde{M}_{t+1}^{(4)}[k_0] &\geq 0.\end{aligned}$$

Assumption 5.3. Preference parameters θ_k satisfy the following inequalities for $k \geq 1$

$$\theta_k \leq 0 \text{ if } k \text{ is even and } \theta_k \geq 0 \text{ if } k \text{ is odd}$$

Assumption 5.3 needs to be strengthened as follows to obtain the completely nonparametric bound in Corollary 5.6:

Assumption 5.4. The first three preference parameters can be expressed as

$$\theta_k = \frac{(-1)^{k+1}}{R_{f,t+1}^k} \text{ for } k \in \{1, 2, 3\}.$$

Remark. Chabi-Yo and Loudis (2020) discuss the economic relevance of these assumptions. Assumption 5.2 concerns odd moments of excess market returns, which are typically negative, since they relate to unfavorable market conditions. Assumption 5.3 is natural given Assumption 5.2, since investors require compensation for exposure to risk-neutral moments. Assumption 5.4 strengthens Assumption 5.3 and is needed to obtain a completely nonparametric bound in Corollary 5.6. One can test the validity of Assumption 5.4 in the data. Chabi-Yo and Loudis (2020) do so and find that Assumption 5.4 cannot be rejected.

Under these assumptions, we can bound the discrepancy between the conditional physical and risk-neutral distribution.

Theorem 5.5 (Lower bound). *Let assumptions 5.2 and 5.3 hold. Assume that the risk-neutral CDF is absolutely continuous w.r.t. Lebesgue measure and $\sup_k \|R_{m,t+1}\|_k := \sup_k \tilde{\mathbb{E}}(|R_{m,t+1}|^k)^{1/k} < \infty$. Finally, define τ^* so that*

$$\tilde{Q}_{t,\tau^*} = R_{f,t+1} - \sup_k \|R_{m,t+1} - R_{f,t+1}\|_k.$$

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

$$\tau - \mathbb{P}_t \left(R_{m,t+1} \leq \tilde{Q}_{t,\tau} \right) \geq \left(\frac{\sum_{k=1}^3 \theta_k \left(\tau \tilde{\mathbb{M}}_{t+1}^{(k)} - \tilde{\mathbb{M}}_{t+1}^{(k)}[\tilde{Q}_{t,\tau}] \right)}{1 + \sum_{k=1}^3 \theta_k \tilde{\mathbb{M}}_{t+1}^{(k)}} \right).$$

Proof. See Appendix A.5. ■

Corollary 5.6. *If, additionally, Assumption 5.4 holds, then for all $\tau \leq \tau^*$*

$$\tau - \mathbb{P}_t \left(R_{m,t+1} \leq \tilde{Q}_{t,\tau} \right) \geq \left(\frac{\sum_{k=1}^3 \frac{(-1)^{k+1}}{R_{f,t+1}^k} \left(\tau \tilde{\mathbb{M}}_{t+1}^{(k)} - \tilde{\mathbb{M}}_{t+1}^{(k)}[\tilde{Q}_{t,\tau}] \right)}{1 + \sum_{k=1}^3 \frac{(-1)^{k+1}}{R_{f,t+1}^k} \tilde{\mathbb{M}}_{t+1}^{(k)}} \right) =: LRB_t(\tau). \quad (5.9)$$

Proof. See Appendix A.5. ■

Under the working hypothesis that the remainder term in the quantile approximation in (5.4) is negligible, the following Corollary is immediate.

Corollary 5.7. *Suppose that the remainder term in (5.4) is negligible, and Assumptions 5.2–5.4 hold, then for all $\tau \leq \tau^*$*

$$Q_{t,\tau} - \tilde{Q}_{t,\tau} \geq \frac{LRB_t(\tau)}{\tilde{f}_t(\tilde{Q}_{t,\tau})}. \quad (5.10)$$

Theorem 5.5 establishes a lower bound on how far the risk-neutral distribution can diverge from the physical distribution conditional on time t . The lower bound that results depends on unknown preference parameters θ_k , which, in principle, can be estimated from return and option data (Chabi-Yo and Loudis, 2020, Section 3.3). The bound in Corollary 5.6 does not require any estimation and can be calculated solely based on time t information. Corollary 5.6 complements the recent literature on the recovery of beliefs. Ross (2015) shows that one can recover $F_t(\cdot)$, if the pricing kernel is transition independent. Subsequent work (Borovička et al., 2016; Qin et al., 2018; Jackwerth and Menner, 2020) casts doubt on the transition independence assumption and shows that recovery is generally impossible. Complimentary to these results, Corollary 5.6 shows that one can still establish a lower bound in the left tail of the distribution, under a different set of (mild) economic constraints. I will show in Section

5.5 that it's much easier to recover the right tail of the distribution, due to the near absence of risk-adjustment.

Corollary 5.7 establishes an interesting relation between the physical and risk-neutral quantile, showing that the difference between the two can be bounded using option data available at time t . The hypothesis that the remainder term in (5.4) is small is confirmed with simulation results in Appendix E, using the Black and Scholes (1973) model. If the lower bound in (5.10) happens to be tight, one can use it as a predictor variable to forecast the physical quantile (also known as Value-at-Risk). Moreover, this predictor variable is not subject to the historical sample bias critique of Goyal and Welch (2008). I now show that the lower bound in Corollary 5.7 is indeed tight.

5.2 Testing Corollary 5.7 in the data

I first outline the procedure to calculate $\tilde{Q}_{t,\tau}$, $LRB_t(\tau)$ and $\tilde{f}_t(\tilde{Q}_{t,\tau})$, which are needed to compute the lower bound in Corollary 5.7. I use the same option data and returns on the S&P500 index, which were used to estimate the quantile bound in Section 4. However, this time we use the full set of daily observations, instead of only using observations at the middle of the month. On every day t , I use the procedure outlined in Appendix B.2 to estimate $\tilde{Q}_{t,\tau}$. Additionally, since $\frac{d}{d\tau}\tilde{Q}_t(\tau) = 1/\tilde{f}_t(\tilde{Q}_{t,\tau})$, I approximate the denominator term in (5.10) by

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

where h is the bandwidth of the τ -grid. Finally, to calculate $LRB_t(\tau)$, I use the estimated quantile curve $\tilde{Q}_{t,\tau}$ in combination with the formula for higher order risk-neutral moments in Appendix A.4.

I use 30/60 day time horizons for testing, starting from January 2, 2004 until December/November 2019. This yields a total of 3361 and 3956 time observations for the 30 and 60 day returns respectively.⁷ Since virtually all macro-finance models assume that the market return is negatively correlated with the SDF, one expects the market return to have a higher probability of

⁷The number of observations is different, since I discard all days where I cannot estimate the risk-neutral curve, for example due to an insufficient number of option data. This happens more often for 30 day returns.

a crash under risk-neutral measure than under physical measure. This means that Corollary 5.7 has non trivial content if $LRB_t(\tau) \geq 0$ in the data. I confirm that this is the case for all dates considered, using the same τ 's from Table 3 below.

5.2.1 In-sample performance

To test whether the lower bound in Corollary 5.7 is tight, I use the following estimator for the latent conditional quantile function

$$\widehat{Q}_{t,\tau} = \widetilde{Q}_{t,\tau} + \underbrace{\frac{LRB_t(\tau)}{\widetilde{f}_t(\widetilde{Q}_{t,\tau})}}_{\text{risk adjustment}}. \quad (5.11)$$

This can be computed at the start of each time period t , using the option data mentioned above. Subsequently, I estimate the coefficients $[\beta_0(\tau), \beta_1(\tau)]^\top$ in the model

$$Q_{t,\tau}(R_{m,t+1}) = \beta_0(\tau) + \beta_1(\tau)\widehat{Q}_{t,\tau}, \quad (5.12)$$

using quantile regression (Koenker and Bassett, 1978). If (5.11) is a good predictor of the conditional quantile function, we expect

$$H_0 : \quad \beta_0(\tau) = 0, \quad \beta_1(\tau) = 1. \quad (5.13)$$

Table 3 summarizes the estimates of model (5.12) for several quantiles. The results are quite striking, as the point estimates of $[\beta_0(\tau), \beta_1(\tau)]^\top$ are close to the $[0, 1]^\top$ benchmark. Moreover, the joint restriction in (5.13) is not rejected for all days and quantiles, thus lending support for a tight lower bound in Corollary 5.7.

The standard errors in Table 3 are obtained via the smooth extended tapered block bootstrap (SETBB) proposed by Gregory et al. (2018). We have to modify the typical standard error estimates in quantile regression, due to the use of overlapping returns as dependent variable which creates serial correlation in the error term. This overlapping data problem has been addressed before in the literature (Hansen and Hodrick, 1980; Hodrick, 1992), but these methods are tailored to OLS and unsuited for quantile regression. The SETBB renders

correct standard errors if the data are weakly dependent, as well as an estimate of the covariance matrix of $[\hat{\beta}_0(\tau), \hat{\beta}_1(\tau)]^\top$. Let $\widehat{\Sigma}(\tau)$ denote the estimated covariance matrix and write $\widehat{\theta}(\tau) := [\hat{\beta}_0(\tau), \hat{\beta}_1(\tau)]^\top - [0, 1]^\top$, then we can test the joint restriction in (5.13) using the Wald statistic

$$T \left[\widehat{\theta}(\tau) \widehat{\Sigma}(\tau)^{-1} \widehat{\theta}(\tau) \right] \rightsquigarrow \chi_2^2, \quad (5.14)$$

where χ_2^2 is the χ -squared distribution with 2 degrees of freedom. I use the prediction horizon as the block length in the bootstrap procedure, which is the only user required input for SETBB.⁸

Table 3: Quantile regression estimates of (5.12)

Maturity:	30 days					60 days				
	$\hat{\beta}_0(\tau)$	$\hat{\beta}_1(\tau)$	Wald test	$R^1(\tau)$ [%]	$R_{\text{OOS}}^1(\tau)$ [%]	$\hat{\beta}_0(\tau)$	$\hat{\beta}_1(\tau)$	Wald test	$R^1(\tau)$ [%]	$R_{\text{OOS}}^1(\tau)$ [%]
$\tau = 0.01$	0.06 (0.3132)	0.97 (0.3506)	0.97	21.08	17.26	-0.23 (0.3368)	1.36 (0.3927)	0.27	16.96	36.59
$\tau = 0.05$	0.20 (0.2944)	0.80 (0.3130)	0.41	9.27	9.21	-0.02 (0.3881)	1.05 (0.4251)	0.98	8.06	22.04
$\tau = 0.1$	0.17 (0.2661)	0.83 (0.2766)	0.46	5.67	6.14	0.13 (0.4056)	0.87 (0.4292)	0.82	5.04	13.96
$\tau = 0.2$	0.21 (0.3808)	0.79 (0.3881)	0.54	1.7	3.78	0.07 (0.4481)	0.93 (0.4581)	0.95	1.57	5.94

Note: Standard errors are shown in parentheses and calculated using the SETBB method of Gregory et al. (2018) with block length equal to the maturity length and 1,000 Monte Carlo bootstrap samples. 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 (5.16). $R_{\text{OOS}}^1(\tau)$ is the out-of-sample goodness-of fit (5.17), using a rolling window size of $10 \times$ maturity.

5.2.2 Out-of-sample performance

Since the in-sample results suggest $\beta_0(\tau) = 0$ and $\beta_1(\tau) = 1$, it is natural to test how well this works out-of-sample by evaluating the predictive model

$$Q_{t,\tau}(R_{m,t+1}) = \widehat{Q}_{t,\tau}.$$

It is worth emphasizing that the left hand side is unknown at time t , whereas the right hand side is known at time t . Moreover, this predictive model does not require any estimation and parallels the discussion in Martin (2017), who used the same idea to measure the conditional equity premium.

An initial concern about the out-of-sample predictor (5.11) is the possibility of crossing, which means that the predicted quantiles $\widehat{Q}_{t,\tau}$ are not monotone with respect to τ . This problem frequently arises in dynamic quantile models

⁸To calculate $\widehat{\Sigma}(\tau)$ with SETBB, I use the `qregBB` function from the `R` package `qregBB` available on the author's Github page: <https://rdrr.io/github/gregorkb/qregBB/man/qregBB.html>.

(Gouriéroux and Jasiak, 2008). It appears, however, not a concern for the predictor in (5.11), as crossing does not occur for either prediction horizon, using the τ 's from Table 3.

To assess the out-of-sample performance further, we like to use a single evaluation metric that summarizes the gains over a simple benchmark. Campbell and Thompson (2008) propose the following out-of-sample R^2 which is popular to evaluate models that predict the equity premium

$$R_{oos}^2 = 1 - \frac{\sum \nu_t^2}{\sum e_t^2}. \quad (5.15)$$

Here, ν_t is the forecast error using a regression model and e_t is the forecast error using a historical (rolling) mean as a predictor, which serves as the benchmark.

However, we cannot use R_{oos}^2 to evaluate the quantile predictor in (5.11), since this out-of-sample metric is appropriate for OLS regression only. Despite this, a natural substitute for quantile regression is available. Koenker and Machado (1999) proposed a goodness-of fit criterion for quantile regression, which resembles the traditional OLS R^2 . In our context, the criterion is given by

$$R^1(\tau) = 1 - \frac{\min_{b_0, b_1} \sum \rho_\tau(R_{m,t+1} - b_0 - b_1 \hat{Q}_{t,\tau})}{\min_{b_0} \sum \rho_\tau(R_{m,t+1} - b_0)}, \quad (5.16)$$

where $\rho_\tau(x) = x(\tau - \mathbb{1}(x < 0))$ is the loss function from quantile regression (Koenker and Bassett, 1978) and $\hat{Q}_{t,\tau}$ is the predictor from (5.11). It is well known that b_0 in the denominator of (5.16) equals the in-sample τ -quantile. The natural out-of-sample analogue of (5.15) for quantile regression is then

$$R_{oos}^1(\tau) := 1 - \frac{\sum \rho_\tau(R_{m,t+1} - \hat{Q}_{t,\tau})}{\sum \rho_\tau(R_{m,t+1} - \bar{Q}_{t,\tau})}, \quad (5.17)$$

where $\bar{Q}_{t,\tau}$ is the historical rolling quantile of the market return until time t . The out-of-sample $R_{oos}^1(\tau)$ is also displayed in Table 3. The predictor variable $\hat{Q}_{t,\tau}$ improves upon the historical rolling quantile out-of-sample in all cases. This is particularly true in the tail of the distribution, 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).

5.3 Time varying disaster risk and dark matter

The in- and out-of-sample results show that $\widehat{Q}_{t,\tau}$ is a good proxy for the latent conditional quantile function. The time fluctuation in $\widehat{Q}_{t,\tau}$ for small τ can therefore be interpreted as time varying disaster risk. The left panels of Figure 5 show the evolution of $\widehat{Q}_{t,\tau}(R_{m,t+1})$ over time for the 30 and 60 day horizon, with $\tau = 0.05$. The time fluctuation in both series is evident from the graph and provides empirical evidence for the thesis that disaster risk is time-varying, as in the models of Gabaix (2012) and Wachter (2013).

For the 30 and 60 day horizon, the lowest quantile forecasts are produced over the 4th quarter of 2008, when the great recession reached its peak. Hence, outlooks were particularly grim during this period and the minimum 30 day 0.05-quantile forecast dropped to 72%, suggesting that a loss of 28% or more had an expected probability of 5%. To put things in perspective, it happened only once since 1926 that the S&P500 index recorded a monthly loss of 28% or more.⁹ Hence, based on historical estimates, a back of the envelope calculation puts the probability of a loss of 28% or more at 1/1139, which is 57 times lower than 5%.

Ross (2015) refers to the impact that changes in perceived disaster probabilities can have on asset prices as *dark matter*: “It is unseen and not directly observable but it exerts a force that can change over time and that can profoundly influence markets”. We can illuminate this dark matter somewhat, since our proxy for $Q_{t,\tau}$ in Figure 5 can be interpreted as a measure of perceived disaster risk. In particular, the calculations above show that this perceived disaster risk differs markedly from historical estimates. Moreover, we observe the most severe downward spikes in Figure 5 during a period associated with large negative returns on the stock market. This dependence points to a kind of leverage effect¹⁰; the negative association between current stock returns and the (perceived) forward looking disaster probability. It would be interesting to incorporate this leverage effect explicitly in a model that features time varying disaster risk and analyze its impact on the equity premium (e.g. in the model

⁹I use historical monthly SP500 return from WRDS, which are available from January 1926 and renders a total of 1139 observations.

¹⁰In finance, the leverage effect typically refers to the negative dependence between returns and changes in volatility (see, e.g. Ait-Sahalia et al. (2013)).

of Wachter (2013)). I leave this for future research.

The right panels of Figure 5 give another interesting view on this dark matter, as they show the evolution of $Q_{t,\tau} - \tilde{Q}_{t,\tau}$ over time. The largest spikes occur once more at the height of the financial crisis and the difference between the two can be as large as 8% (30 day horizon) or 15% (60 day horizon). This suggests that the risk-neutral quantile decreases disproportionately more than the physical quantile during crises.

It is well known that risk-neutral quantiles are less than historical quantiles in the left tail of the distribution. As Ross (2015) remarks, this means that either the market forecasts a higher probability of a loss than historical estimates suggest, or the market requires a high risk premium for insurance against crash risk. The former is a statement about the physical distribution, whereas the latter concerns the risk-neutral distribution. Since market forecasts of a crash are unobserved one cannot separate the two effects. However, our approach allows us to separate the effects, because we have a good proxy for $Q_{t,\tau}$. Figure 5 suggests that the insurance effect is more dominant than the market expectation effect during a crisis. Namely, during the great recession, market expectations of the severity of a crash went down (left panels of Figure 5), but the risk-neutral quantile decreased even more (right panels of Figure 5). In other words, risk aversion changes more than the perceived disaster risk.

5.4 Lognormal returns and risk-neutral quantiles

The previous Section compared the out-of-sample performance of $\hat{Q}_{t,\tau}$ to the (rolling) historical quantile benchmark. There is another natural benchmark to consider, namely the risk-neutral quantile $\tilde{Q}_{t,\tau}$. Since $\tilde{Q}_{t,\tau}$ is risk adjusted it may not seem a good idea to use it as a predictor. This is why we introduced the risk adjustment term $LRB_t(\tau)/\tilde{f}_t(\tilde{Q}_{t,\tau})$ in the first place. Somewhat surprisingly, I now show that among the class of lognormal models, $\tilde{Q}_{t,\tau}$ is the optimal predictor of the physical quantile, due to rotation symmetry of the quantile regression estimator.

To describe the environment, consider the following discretized version of the Black-Scholes model, with a riskless asset that offers a certain return $R_f = e^{r_f\lambda}$ and a risky asset with return $R_{t+1} = \exp([\mu - \frac{1}{2}\sigma_t^2]\lambda + \sigma_t\sqrt{\lambda}Z_{t+1})$, where Z_{t+1}

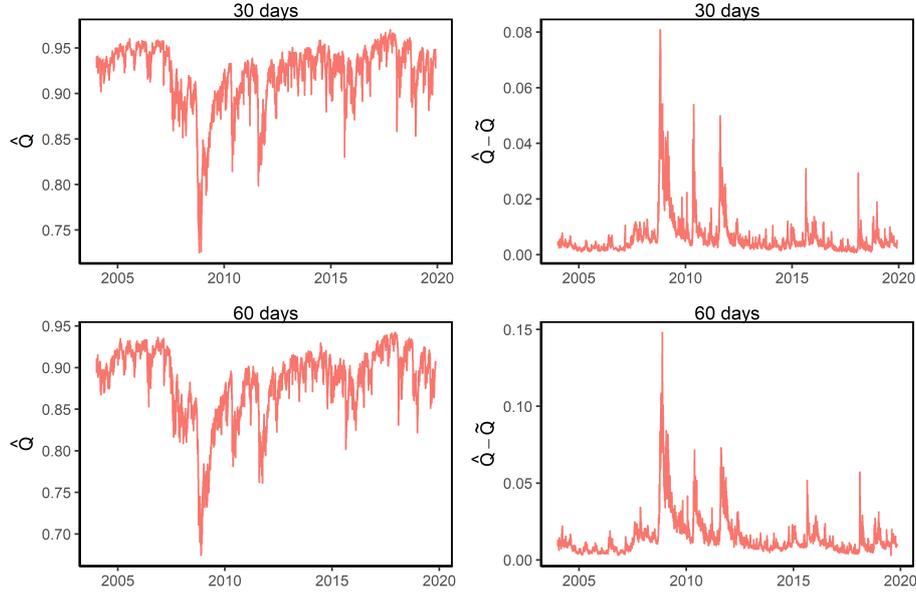


Figure 5: The left panels show the real time quantile predictor $\hat{Q}_{t,\tau}$ from (5.11) for $\tau = 0.05$, using prediction horizons 30 and 60 days. The right panels show the difference between $\hat{Q}_{t,\tau} - \tilde{Q}_{t,\tau}$ for $\tau = 0.05$.

is standard normal, σ_t is the conditional (\mathcal{F}_t -measurable) volatility of returns and λ denotes the time difference in years between period t and $t + 1$. In this setup, $M_{t+1} := \exp(-[r_f + \xi_t^2/2]\lambda - \xi_t\sqrt{\lambda}Z_{t+1})$, is a valid SDF with conditional Sharpe ratio

$$\xi_t = \frac{\mu - r_f}{\sigma_t}.$$

The implied dynamics under risk-neutral measure are given by

$$R_{t+1} = \exp\left(\left(r_f - \frac{1}{2}\sigma_t^2\right)\lambda + \sigma_t\sqrt{\lambda}Z_{t+1}\right). \quad (5.18)$$

The next Theorem shows that quantile regression using the risk-neutral quantile as a regressor renders an identical forecast when using the (unobserved) physical quantile as a regressor.

Theorem 5.8. *Consider the lognormal model described above with return observations $\{R_{t+1}\}_{t=1}^T$ stacked in the $T \times 1$ vector R . Let $\tilde{X}_t(\tau) := [1 \ \tilde{Q}_{t,\tau}(R_{t+1})]^\top$ and denote the $T \times 2$ matrix of stacked $\tilde{X}_t(\tau)$ by $\tilde{X}(\tau)$. Define the regression quantile $\hat{\beta}(\tau; R, \tilde{X}(\tau))$ as the solution to the quantile regression with the risk-*

neutral quantile as a covariate

$$\widehat{\beta}(\tau; R, \widetilde{X}(\tau)) \in \arg \min_{\beta \in \mathbb{R}^2} \sum_{t=1}^T \rho_{\tau} \left(R_{t+1} - \widetilde{X}_t(\tau)^{\top} \beta \right),$$

where $\rho_{\tau}(u) = (\tau - \mathbb{1}(u \leq 0))u$.

Similarly, let $X_t(\tau) := [1 \ Q_{t,\tau}(R_{t+1})]^{\top}$, $X(\tau)$ the $T \times 2$ matrix of stacked $X_t(\tau)$ and define $\widehat{\beta}(\tau; R, X(\tau))$ as the solution to the quantile regression using the physical quantile as a covariate

$$\widehat{\beta}(\tau; R, X(\tau)) \in \arg \min_{\beta \in \mathbb{R}^2} \sum_{t=1}^T \rho_{\tau} \left(R_{t+1} - X_t(\tau)^{\top} \beta \right). \quad (5.19)$$

Then

$$\widetilde{X}_{T+1}(\tau)^{\top} \widehat{\beta}(\tau; R, \widetilde{X}(\tau)) = X_{T+1}(\tau)^{\top} \widehat{\beta}(\tau; R, X(\tau)). \quad (5.20)$$

Proof. See Appendix A.6. ■

In sum, if returns are conditionally lognormal with time varying volatility there is no need to risk adjust, since the quantile forecast based on the risk-neutral quantile satisfies

$$Q_{T+1}(\tau) = \widetilde{Q}_{T+1}(\tau) \widehat{\beta}_1(\tau) + o_p(1) \quad \text{as } T \rightarrow \infty.$$

Intuitively, the reason is that $\widehat{\beta}_1(\tau)$ picks up the risk premium, so that $\widetilde{Q}_{t,\tau} \widehat{\beta}_1(\tau)$ rotates back into physical quantile units. A similar situation occurs in Principal Component Analysis, where it is enough to identify the principal component up to some rotation matrix to make predictions (Bai, 2003).

The assumption underlying the result is that the only source of variation in the distribution of returns is changes in conditional volatility. This is in essence the same idea underlying the popular GARCH models. However, I have abstracted away from specifying what actually drives the volatility process, as opposed to GARCH type models. Hence, the result of Theorem 5.8 is valid for any conditional volatility specification. The result comes at the cost of modeling the returns as conditionally lognormal. There is ample evidence that returns are not conditionally lognormal (Martin, 2017), but given the popularity of the

lognormal assumption in financial models it is still an interesting benchmark to consider.

Theorem 5.8 can actually be used as model free evidence against the conditional lognormal assumption. To see this, we can use quantile regression to estimate coefficients $\widehat{\beta}_{0,t}(\tau) + \widehat{\beta}_{1,t}(\tau)$, using $\widetilde{Q}_{t,\tau}$ as the only covariate to explain the quantiles of $R_{m,t+1}$. The t -subscript in $\beta_{.,t}$ refers to the fact that the coefficients are estimated using information up to time t . I use an expanding window to estimate $\beta_{.,t}$ based on quantile regression and then produce dynamic quantile forecasts of the form

$$\widehat{Q}_{t,\tau} = \widehat{\beta}_{0,t}(\tau) + \widehat{\beta}_{1,t}(\tau)\widetilde{Q}_{t,\tau}. \quad (5.21)$$

If returns follow the lognormal dynamics in (5.18), Theorem 5.8 suggests that

$$\widehat{Q}_{t,\tau} \approx Q_{t,\tau}.$$

This hypothesis can be formalized by testing for the joint restriction

$$H_0 : [\beta_0(\tau), \beta_1(\tau)] = [0, 1]^\top, \quad (5.22)$$

in the quantile regression

$$\min_{\beta_0, \beta_1} \sum_t \rho_\tau \left(R_{m,t+1} - \beta_0 - \beta_1 \widehat{Q}_{t,\tau} \right).$$

The same Wald restriction test is used as in (5.14), using SETBB to compute the covariance matrix. The results are summarized in Table 4. The Wald test on the joint restriction is rejected for any τ and prediction horizon. Moreover, the explanatory power is low, as $R^1(\tau)$ is at most a couple of percentage points. This is strong evidence against the conditional lognormal assumption. The inconsistencies that arise from the conditional lognormal assumption have been documented before by Martin (2017, Result 4) using a different method, exploiting the difference between the VIX and SVIX index. The conclusion from that paper applies more generally though, since it allows for time variation in the mean and risk-free rate.

Table 4: Quantile regression with risk-neutral covariate

Horizon	30 days				60 days			
	$\hat{\beta}_0(\tau)$	$\hat{\beta}_1(\tau)$	Wald test	$R^1(\tau)[\%]$	$\hat{\beta}_0(\tau)$	$\hat{\beta}_1(\tau)$	Wald test	$R^1(\tau)[\%]$
$\tau = 0.01$	0.62 (0.1912)	0.30 (0.2066)	0	5.43	0.75 (0.1603)	-0.01 (0.2287)	0	0.03
$\tau = 0.05$	0.61 (0.1810)	0.35 (0.1894)	0	6.52	0.69 (0.2687)	0.24 (0.2828)	0	2.33
$\tau = 0.1$	0.59 (0.1860)	0.38 (0.1927)	0	4.02	0.71 (0.2978)	0.24 (0.3103)	0	1.49
$\tau = 0.2$	0.54 (0.2412)	0.45 (0.2461)	0	1.44	0.71 (0.3276)	0.27 (0.3344)	0	1.02
Observations	2832				3397			

Note: Quantile regression of (5.21) using an expanding window based on an initial 500 observations. Wald test denotes the p -value of the joint restriction $[\beta_0(\tau), \beta_1(\tau)]^\top = [0, 1]^\top$.

5.5 Asymmetric tail behavior

The disaster risk specification of Backus et al. (2011), or the time varying disaster risk model of Wachter (2013) imply another testable implication that is not shared with the LRR model. Namely, shocks to consumption are assumed to be negative should a disaster occur. This implies the fat left tail in the risk-neutral distribution shown in Figure 2, but the right tail of the risk-neutral and physical distribution are nearly identical. In contrast, the LRR model implies that the risk-neutral and physical distribution differ most around the median due to the conditional lognormal assumption (see Figure 6).

These qualitatively different hypotheses can be tested by running the quantile regression

$$Q_{t,\tau} = \beta_0(\tau) + \beta_1(\tau)\tilde{Q}_{t,\tau}, \quad (5.23)$$

and testing the null hypothesis $[\beta_0(\tau), \beta_1(\tau)]^\top = [0, 1]^\top$. For example, the disaster risk model implies rejection of H_0 for small τ , but non-rejection for τ close to one, since the risk-neutral quantile is almost identical to the physical quantile in the right tail. The LRR model implies rejection around $\tau = 0.5$ and non-rejection for τ close to 0 or 1.

Table 5 contains the result. Consistent with the disaster risk model, we see quite different behavior in the left and right tail. The regression estimates for small τ are quite different from the $[0, 1]^\top$ benchmark and the null hypothesis is rejected or marginally not rejected at 5% significance. In contrast, the regression estimates for τ closer to 1 indicate that $\tilde{Q}_{t,\tau}$ is a good predictor of the latent quantile function $Q_{t,\tau}$. The p -values for the joint restriction are rather high, even

for the median. Following [Engle and Manganelli \(2004\)](#), [Table 5](#) also reports the expected value of the hit function, defined as

$$\text{Hit}_t = \mathbb{1} \left(R_{m,t+1} < \tilde{Q}_{t,\tau} \right) - \tau. \quad (5.24)$$

The expected value is 0 if $Q_{t,\tau} = \tilde{Q}_{t,\tau}$, that is, if investors are risk-neutral. Again consistent with the disaster risk, the expected value of [\(5.24\)](#) is more negative in the left than right tail. Somewhat surprisingly, the expected value is positive for $\tau = 0.95$. Such a finding cannot be warranted by a monotonically decreasing pricing kernel, but is consistent with a so called *U-shaped* kernel, as in [Bakshi et al. \(2010\)](#).¹¹

[Table 5](#) also connects to the recovery of beliefs ([Ross, 2015](#); [Borovička et al., 2016](#)). Namely, the right tail of the physical distribution can be accurately recovered from the right tail of the risk-neutral distribution, since there is virtually no risk adjustment in that part of the distribution. This observation may only have limited use in practice though, since it is the left tail of both distributions that is typically of interest in economic applications, such as Value-at-Risk.

Table 5: Risk-neutral prediction

Horizon	30 days				60 days			
	$\hat{\beta}_0(\tau)$	$\hat{\beta}_1(\tau)$	Wald test	Hit	$\hat{\beta}_0(\tau)$	$\hat{\beta}_1(\tau)$	Wald test	Hit
$\tau = 0.05$	0.31 (0.2510)	0.69 (0.2680)	0.06	-0.03	0.22 (0.3293)	0.79 (0.3639)	0.46	-0.03
$\tau = 0.1$	0.32 (0.2273)	0.67 (0.2372)	0.02	-0.04	0.38 (0.3074)	0.62 (0.3279)	0.06	-0.05
$\tau = 0.2$	0.38 (0.3211)	0.62 (0.3278)	0.07	-0.04	0.49 (0.3266)	0.51 (0.3359)	0.01	-0.06
$\tau = 0.5$	0.06 (0.2273)	0.94 (0.2258)	0.88	-0.07	0.13 (0.2982)	0.87 (0.2946)	0.69	-0.04
$\tau = 0.8$	-0.04 (0.1842)	1.04 (0.1788)	0.92	-0.02	0.04 (0.2377)	0.96 (0.2270)	0.95	-0.02
$\tau = 0.9$	0.04 (0.1547)	0.96 (0.1486)	0.88	0.01	0.08 (0.1951)	0.93 (0.1837)	0.72	0
$\tau = 0.95$	0.00 (0.1518)	1.00 (0.1445)	1	0.01	0.04 (0.1682)	0.96 (0.1565)	0.88	0.02
Observations	3361				3956			

Note: *Quantile regression of (5.23)*. Standard errors are based on SETBB and Wald test denotes the *p-value of the joint restriction test* $[\beta_0(\tau), \beta_1(\tau)]^T = [0, 1]^T$.

¹¹It would require a formal test to argue that the expected value Hit_t is statistically different from 0. I leave this for future research.

6 Conclusion

This paper proposes a new bound (quantile bound) on the SDF volatility, which, in contrast to the HJ bound, compares the physical and risk-neutral distribution at every τ -quantile. I show that the quantile bound compares favorably to the HJ bound in scenarios where returns are heavy tailed or models that incorporate disaster risk. Among others, the quantile bound uses information of the data beyond the mean and variance, which are central to the HJ bound. In the data, I find suggestive evidence of the presence of disaster risk and show that the quantile bound is stronger than the HJ bound. I argue that this points to misspecification of asset pricing models, such as CAPM or LRR, since they fail to incorporate disaster risk.

Subsequently, I analyze the conditional difference between physical and risk-neutral quantiles. A von Mises expansion is used to analyze this difference, which, under mild economic constraints, can be bounded by moments of the risk-neutral distribution. Quantile regression estimates confirm that the lower bound predicts well in the data and that the lower bound is tight. This finding complements the vast literature in finance which is concerned with conditional mean forecasts.

I develop the analogy with the conditional mean further and propose a measure of out-of-sample performance, which shows that the lower bound outperforms the historical and risk-neutral quantile in the left tail of the distribution. This evidence leads me to interpret the lower bound as a good approximation of the latent conditional quantile function. I find that the lower bound fluctuates significantly over time, pointing towards time changing disaster risk. In addition, the lower bound drops markedly during periods associated with market distress, indicating that perceived disaster risk is correlated with contemporaneous disasters. I then compare the difference between the lower bound and risk-neutral quantile and argue that, during the great recession, an increase in the premium for insurance against disaster risk was more pronounced than the increase in perceived disaster risk. This is a novel way to disentangle two effects which are normally indistinguishable.

A final application considers the conditional implications of asset pricing models. I establish that the risk-neutral quantile severely underestimates the

physical quantile in the left tail, but not in the right tail. This is in line with the class of disaster risk models, but contradicts the LRR model. Overall, the results underscore the importance of considering features of the data beyond the mean and variance and provide deeper insight in the influence of tail risk on asset prices.

A Proofs and computations

A.1 Proof of Theorem 2.1 (quantile bound)

Proof. I suppress the dependence of the τ -quantile on R and write $\tilde{Q}_\tau := \tilde{Q}_\tau(R)$. Start from the definition of a risk-neutral quantile

$$\begin{aligned}\tau &= \tilde{\mathbb{P}}[R \leq \tilde{Q}_\tau] = \tilde{\mathbb{E}}[\mathbb{1}(R \leq \tilde{Q}_\tau)] = R_f \mathbb{E}_t[M \mathbb{1}(R \leq \tilde{Q}_\tau)] \\ &= R_f [\text{COV}(M, \mathbb{1}(R \leq \tilde{Q}_\tau)) + \mathbb{E}[M] \mathbb{E}[\mathbb{1}(R \leq \tilde{Q}_\tau)]] \\ &= R_f \underbrace{\text{COV}(M, \mathbb{1}(R \leq \tilde{Q}_\tau))}_{=\mathbb{P}(R \leq \tilde{Q}_\tau)} + \mathbb{E}[\mathbb{1}(R \leq \tilde{Q}_\tau)].\end{aligned}\tag{A.1}$$

Rearranging then yields

$$\frac{\tau - \mathbb{P}(R \leq \tilde{Q}_\tau)}{R_f} = \text{COV}(M, \mathbb{1}(R \leq \tilde{Q}_\tau)).$$

Using Cauchy-Schwarz renders the inequality

$$\begin{aligned}\left| \frac{\tau - \mathbb{P}(R \leq \tilde{Q}_\tau)}{R_f} \right| &\leq \sigma(M) \sigma(\mathbb{1}(R \leq \tilde{Q}_\tau)) \\ \frac{|\tau - \mathbb{P}(R \leq \tilde{Q}_\tau)|}{\sigma(\mathbb{1}(R \leq \tilde{Q}_\tau)) R_f} &\leq \sigma(M).\end{aligned}\tag{A.2}$$

Finally, since $\mathbb{1}(R \leq \tilde{Q}_\tau)$ is a Bernoulli random variable, it follows that

$$\sigma(\mathbb{1}(R \leq \tilde{Q}_\tau)) = \sqrt{\mathbb{P}(\tilde{Q}_\tau(R)) \times (1 - \mathbb{P}(\tilde{Q}_\tau(R)))}.\tag{A.3}$$

Theorem 2.1 now follows after substituting (A.3) into (A.2). ■

A.2 Quantile bound Pareto distribution

This Section proves the two propositions in the main text about the quantile bound, when the distribution of returns is Pareto.

Proof of Proposition 2.2. (i) The distribution of returns is Pareto, since

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

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

$$\mathbb{E}[R] = \frac{B}{1-\beta} \quad \sigma^2(R) = \frac{B^2}{1-2\beta} - \left(\frac{B}{1-\beta}\right)^2. \quad (\text{A.4})$$

Likewise, the distribution of the SDF follows from

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

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 $\mathbb{E}[MR] = 1$ forces

$$\frac{AB}{\alpha - \beta + 1} = 1. \quad (\text{A.5})$$

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

$$\frac{A}{\alpha + 1} = \frac{1}{R_f}. \quad (\text{A.6})$$

The Sharpe ratio can now be computed from (A.4) and (A.6). ■

Proof of Proposition 2.3. (i) Since $R_f M$ is the Radon-Nikodym derivative that

induces a change of measure from \mathbb{P} to $\tilde{\mathbb{P}}$, it follows that

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

The last line follows from (A.6).

(ii) It is easy to show that the quantiles of a $\mathbf{PAR}(C, \zeta)$ distribution are given by

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

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

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

Hence, the quantile bound evaluates to

$$\frac{|\tau - \mathbb{P}(R \leq \tilde{Q}_\tau)|}{R_f \sigma(\mathbb{1}(R \leq \tilde{Q}_\tau))} = \frac{A}{1 + \alpha} \frac{|\tau - 1 + (1 - \tau)^{\frac{1}{\alpha+1}}|}{\sqrt{(1 - (1 - \tau)^{\frac{1}{\alpha+1}})(1 - \tau)^{\frac{1}{\alpha+1}}}}. \quad (\text{A.7})$$

(iii). The HJ bound, as given by the Sharpe ratio in (2.9), goes to 0 as $\beta \uparrow 1/2$ since $\sigma(R) \uparrow \infty$. \blacksquare

A.3 Lognormal return and SDF

This Section provides a closed form approximation for the relative efficiency between the HJ and quantile bound under joint lognormality. Let

$$\begin{aligned} R_{t+1} &= e^{(\mu_R - \frac{\sigma_R^2}{2})\lambda + \sigma_R\sqrt{\lambda}Z_R} \\ M_{t+1} &= e^{-(r_f + \frac{\sigma_M^2}{2})\lambda + \sigma_M\sqrt{\lambda}Z_M}. \end{aligned}$$

Both Z_R and Z_M are standard normal random variables with correlation ρ . First, approximate M_{t+1} by a first order Taylor expansion, which gives

$$\widehat{M}_{t+1} = e^{-(r_f + \frac{\sigma_M^2}{2})\lambda} + Z_M\sigma_M\sqrt{\lambda}e^{-(r_f + \frac{\sigma_M^2}{2})\lambda}.$$

Notice that $\widehat{M}_{t+1} = M_{t+1} + o_p(\sqrt{\lambda})$. Consequently, by Stein's Lemma

$$\begin{aligned} \text{COV}(R_{t+1}, M_{t+1}) &\approx \text{COV}(R_{t+1}, \widehat{M}_{t+1}) = \sigma_M\sqrt{\lambda}e^{-(r_f + \frac{\sigma_M^2}{2})\lambda} \text{COV}(R_{t+1}, Z_M) \\ &= \sigma_M\sqrt{\lambda}e^{-(r_f + \frac{\sigma_M^2}{2})\lambda} \mathbb{E} \left[\sigma_R\sqrt{\lambda} \exp \left(\left[\mu_R - \frac{\sigma_R^2}{2} \right] \lambda + \sigma_R\sqrt{\lambda}Z_R \right) \right] \text{COV}(Z_R, Z_M) \\ &= \sigma_M\sigma_R\lambda e^{-(r_f + \frac{\sigma_M^2}{2})\lambda} e^{\mu_R\lambda} \text{COV}(Z_R, Z_M). \end{aligned}$$

Again by Stein's Lemma

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

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

$$\left| \frac{\mathbb{E}[R_{t+1}] - e^{\lambda r_f}}{\tau - \mathbb{P}(R_{t+1} \leq \tilde{Q}_\tau)} \right| \approx \frac{\sigma_R\sqrt{\lambda}e^{\mu_R\lambda}}{f(x)}. \quad (\text{A.8})$$

The same reasoning in Example 2.2 implies that the relative efficiency between the HJ and quantile bound can be approximated by

$$\begin{aligned} \frac{\text{HJ bound}}{\text{Quantile bound}} &= \frac{\frac{|\mathbb{E}[R_{t+1}] - R_{f,t+1}|}{\sigma(R_{t+1})R_{f,t+1}}}{\frac{|\tau - \mathbb{P}(R_{t+1} \leq \tilde{Q}_\tau)|}{\sqrt{\mathbb{P}(R_{t+1} \leq \tilde{Q}_\tau)(1 - \mathbb{P}(R_{t+1} \leq \tilde{Q}_\tau))}R_{f,t+1}}} \\ &\stackrel{\text{(A.8)}}{\approx} \frac{\sqrt{\mathbb{P}(r_{t+1} \leq x)(1 - \mathbb{P}(r_{t+1} \leq x))}}{\sigma(R_{t+1})} \times \frac{\sigma_R \sqrt{\lambda} e^{\mu_R \lambda}}{f(x)}. \quad (\text{A.9}) \end{aligned}$$

Here, $r_{t+1} = \log R_{t+1}$ and $x = \log \tilde{Q}_\tau$. Using the same reasoning as in Example 2.2, the expression on the RHS of (A.9) is minimized by choosing $x = \log \tilde{Q}_\tau^*$ s.t. $\mathbb{P}(R_{t+1} \leq \tilde{Q}_\tau^*) = 1/2$. In that case the relative efficiency equals

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

A.4 Formulas for market moments

This Section presents formulas for the (un)truncated risk-neutral moments of excess market return. An alternative way to calculate these is provided in Chabi-Yo and Loudis (2020, Appendix B). I use a slight abuse of notation and write $\tilde{Q}_{R_{t+1}}(\tau) := \tilde{Q}_{t,\tau}(R_{t+1})$, to obviate that the integrals below are taken with respect to τ .

Proposition A.1. *Higher order risk-neutral moments can be computed directly from the risk-neutral quantile function*

$$\tilde{\mathbb{E}}_t [(R_{t+1} - R_{f,t+1})^n] = \int_0^1 [\tilde{Q}_{R_{t+1}-R_{f,t+1}}(\tau)]^n d\tau = \int_0^1 [\tilde{Q}_{R_{t+1}}(\tau) - R_{f,t+1}]^n d\tau. \quad (\text{A.10})$$

And the truncated higher order risk-neutral moments also follow from

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

Where $\tilde{F}_t(x) := \tilde{\mathbb{P}}_t(R_{t+1} \leq x)$ is the risk-neutral CDF. Frequently I use $k_0 = \tilde{Q}_{t,\tau}(R_{t+1})$, in which case the truncated moment formula reduces to

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

Proof. For any random variable X and integer n s.t. the n -th moment exist

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

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.10). The truncated formula follows similarly. \blacksquare

A.5 Proof of Theorem 5.5 and Corollary 5.6

Proof of Theorem 5.5 and Corollary 5.6. I split the proof in three parts.

Part 1: Showing that $\widetilde{\text{COV}}_t \left[\mathbb{1} \left(R_{m,t+1} \leq \tilde{Q}_{t,\tau} \right), (R_{m,t+1} - R_{f,t+1})^k \right] \leq 0$ for k odd.

Temporarily write $X = R_{m,t+1}$. To prove the claim above I distinguish 3 cases. Take two i.i.d. copies X_1, X_2 with the same law as X under risk-neutral measure and consider

$$\Lambda := \underbrace{\left(\mathbb{1} \left(X_1 \leq \tilde{Q}_{t,\tau} \right) - \mathbb{1} \left(X_2 \leq \tilde{Q}_{t,\tau} \right) \right)}_{=I} \underbrace{\left((X_1 - R_{f,t+1})^k - (X_2 - R_{f,t+1})^k \right)}_{=II}. \quad (\text{A.11})$$

Case 1: ($I = -1$). This implies $X_2 < X_1$. Since k is odd I get $II > 0$ so that $\Lambda < 0$.

Case 2: ($I = 0$). This implies $\Lambda = 0$.

Case 3: ($I = 1$). This implies $X_1 \leq X_2$ and hence $II < 0$. Therefore $\Lambda < 0$.

Combining all three cases I get that $\Lambda \leq 0$ almost surely. Take conditional (risk-neutral) expectations on both sides of (A.11), using the non-positivity of Λ and the independence of X_1, X_2 proves that the covariance term is negative. Since by assumption $\theta_k \geq 0$ when k is odd I obtain

$$\begin{aligned} & \tilde{\mathbb{E}}_t \left[\mathbb{1} \left(R_{m,t+1} \leq \tilde{Q}_{t,\tau} \right) (R_{m,t+1} - R_{f,t+1})^k \right] - \tau \tilde{\mathbb{E}}_t \left[(R_{m,t+1} - R_{f,t+1})^k \right] \leq 0 \\ & \implies \tau \tilde{\mathbb{M}}_{t+1}^{(k)} - \tilde{\mathbb{M}}_{t+1}^{(k)}[\tilde{Q}_{t,\tau}] \geq 0 \implies \theta_k \left(\tau \tilde{\mathbb{M}}_{t+1}^{(k)} - \tilde{\mathbb{M}}_{t+1}^{(k)}[\tilde{Q}_{t,\tau}] \right) \geq 0. \end{aligned}$$

Part 2: Showing that $\widetilde{\text{COV}}_t \left[\mathbb{1} \left(R_{m,t+1} \leq \tilde{Q}_{t,\tau} \right), (R_{m,t+1} - R_{f,t+1})^k \right] \geq 0$ for k even and τ small enough.

This requires more delicate reasoning. First note that the covariance term goes to zero as $\tau \rightarrow 0$ as a consequence of the Cauchy-Schwarz inequality and the continuity of probability measures. Hence, to show that for τ small enough the covariance term is positive, it suffices that the covariance term, seen as a function of τ , has positive slope for τ small enough. To show this, write the covariance as

$$\tilde{\mathbb{E}}_t \left[\left(\mathbb{1} \left(R_{m,t+1} \leq \tilde{Q}_{t,\tau} \right) - \tau \right) (R_{m,t+1} - R_{f,t+1})^k \right]. \quad (\text{A.12})$$

Consider the associated function

$$\begin{aligned} \Gamma(\tau) &:= \tilde{\mathbb{E}}_t \left[\left(\mathbb{1} \left(R_{m,t+1} \leq \tilde{Q}_{t,\tau} \right) - \tau \right) (R_{m,t+1} - R_{f,t+1})^k \right] \\ &= \int_{-\infty}^{\tilde{Q}_{t,\tau}} (R - R_{f,t+1})^k \tilde{f}_{R_{m,t+1}}(R) dR - \tau \int_{-\infty}^{\infty} (R - R_{f,t+1})^k \tilde{f}_{R_{m,t+1}}(R) dR. \end{aligned}$$

Here $\tilde{f}_{R_{m,t+1}}(\cdot)$ is the (risk-neutral) PDF of the market return. From Leibniz' rule

$$\begin{aligned} \frac{\partial}{\partial \tau} \Gamma(\tau) &= (\tilde{Q}_{t,\tau} - R_{f,t+1})^k \tilde{f}_{R_{m,t+1}}(\tilde{Q}_{t,\tau}) \frac{\partial \tilde{Q}_{t,\tau}}{\partial \tau} - \tilde{\mathbb{E}}_t \left[(R_{m,t+1} - R_{f,t+1})^k \right] \\ &= (\tilde{Q}_{t,\tau} - R_{f,t+1})^k - \tilde{\mathbb{E}}_t \left[(R_{m,t+1} - R_{f,t+1})^k \right], \end{aligned} \quad (\text{A.13})$$

since, by the rules for derivatives of inverses

$$\frac{\partial \tilde{Q}_{t,\tau}}{\partial \tau} = \frac{1}{\tilde{f}_{R_{m,t+1}}(\tilde{Q}_{t,\tau})}.$$

Because I assume that $\sup_k \|R_{m,t+1}\|_k < \infty$, it follows that (A.13) is positive for all $\tau \in [0, \tau^*]$, where τ^* solves

$$\tilde{Q}_{t,\tau^*} = R_{f,t+1} - \sup_k \|R_{m,t+1} - R_{f,t+1}\|_k.$$

In conclusion, I have shown that the covariance (A.12) vanishes when $\tau \rightarrow 0^+$ and the slope of (A.12) is positive for all $\tau \leq \tau^*$. This means that (A.12) is

positive for all $\tau \leq \tau^*$. Thus for all such $\tau \in (0, \tau^*]$

$$\tau \tilde{\mathbb{M}}_{t+1}^{(k)} - \tilde{\mathbb{M}}_{t+1}^{(k)}[\tilde{Q}_{t,\tau}] \leq 0.$$

Hence, since $\theta_k \leq 0$ for k even

$$\theta_k \left(\tau \tilde{\mathbb{M}}_{t+1}^{(k)} - \tilde{\mathbb{M}}_{t+1}^{(k)}[\tilde{Q}_{t,\tau}] \right) \geq 0.$$

Part 3: Combining both cases

I have now established $\theta_k(\tau \tilde{\mathbb{M}}_{t+1}^{(k)} - \tilde{\mathbb{M}}_{t+1}^{(k)}[\tilde{Q}_{t,\tau}]) \geq 0$ for all k and $\tau \leq \tau^*$.

Therefore

$$\begin{aligned} Q_{t,\tau} - \tilde{Q}_{t,\tau} &\approx \frac{1}{\tilde{f}_t(\tilde{Q}_{t,\tau})} \left(\frac{\sum_{k=1}^{\infty} \theta_k \left(\tau \tilde{\mathbb{M}}_{t+1}^{(k)} - \tilde{\mathbb{M}}_{t+1}^{(k)}[\tilde{Q}_{t,\tau}] \right)}{1 + \sum_{k=1}^{\infty} \theta_k \tilde{\mathbb{M}}_{t+1}^{(k)}} \right) \\ &\geq \frac{1}{\tilde{f}_t(\tilde{Q}_{t,\tau})} \left(\frac{\sum_{k=1}^3 \theta_k \left(\tau \tilde{\mathbb{M}}_{t+1}^{(k)} - \tilde{\mathbb{M}}_{t+1}^{(k)}[\tilde{Q}_{t,\tau}] \right)}{1 + \sum_{k=1}^3 \theta_k \tilde{\mathbb{M}}_{t+1}^{(k)}} \right). \end{aligned} \quad (\text{A.14})$$

If additionally Assumption 5.4 holds, then

$$\theta_1 = \frac{1}{R_{f,t+1}}, \theta_2 = -\frac{1}{R_{f,t+1}^2}, \text{ and } \theta_3 = \frac{1}{R_{f,t+1}^3}.$$

Using this in (A.14) gives

$$Q_{t,\tau} - \tilde{Q}_{t,\tau} \geq \frac{1}{\tilde{f}_t(\tilde{Q}_{t,\tau})} \left(\frac{\sum_{k=1}^3 \frac{(-1)^{k+1}}{R_{f,t+1}^k} \left(\tau \tilde{\mathbb{M}}_{t+1}^{(k)} - \tilde{\mathbb{M}}_{t+1}^{(k)}[\tilde{Q}_{t,\tau}] \right)}{1 + \sum_{k=1}^3 \frac{(-1)^{k+1}}{R_{f,t+1}^k} \tilde{\mathbb{M}}_{t+1}^{(k)}} \right).$$

■

A.6 Proof of Theorem 5.8

Proof. By definition

$$\tau = \mathbb{P}_t(R_{t+1} \leq Q_{t,\tau}) = \mathbb{P}_t \left(\exp \left(-\frac{1}{2} \sigma_t^2 \lambda + \sigma_t \sqrt{\lambda} Z_{t+1} \right) \leq \exp(-\mu \lambda) Q_{t,\tau} \right).$$

Similarly

$$\tau = \tilde{\mathbb{P}}_t(R_{t+1} \leq \tilde{Q}_{t,\tau}) = \tilde{\mathbb{P}}_t \left(\exp \left(-\frac{1}{2} \sigma_t^2 \lambda + \sigma_t \sqrt{\lambda} Z_{t+1} \right) \leq \exp(-r_f \lambda) \tilde{Q}_{t,\tau} \right).$$

As a result

$$e^{(\mu-r_f)\lambda} \tilde{Q}_{t,\tau} = Q_{t,\tau}. \quad (\text{A.15})$$

Recall that the quantile regression estimator is equivariant to reparametrization of design: for any 2×2 nonsingular matrix A , we have

$$\hat{\beta}(\tau; R, XA) = A^{-1} \hat{\beta}(\tau; R, X).$$

By Equation (A.15)

$$X(\tau) = \tilde{X}(\tau) \times \underbrace{\begin{bmatrix} 1 & 0 \\ 0 & e^{(\mu-r_f)\lambda} \end{bmatrix}}_{:=A}.$$

Therefore

$$\hat{\beta}(\tau; R, X(\tau)) = \hat{\beta}(\tau; R, \tilde{X}(\tau)A) = A^{-1} \hat{\beta}(\tau; R, \tilde{X}(\tau)).$$

Hence, the predicted quantile using the physical quantile regression (5.19) equals

$$\begin{aligned} [1 \quad Q_{T+1,\tau}(R_{T+2})] \hat{\beta}(\tau; R, X(\tau)) &= [1 \quad Q_{T+1,\tau}(R_{T+2})] A^{-1} \hat{\beta}(\tau; R, \tilde{X}(\tau)) \\ &= [1 \quad \tilde{Q}_{T+1,\tau}(R_{T+2})] \hat{\beta}(\tau; R, \tilde{X}(\tau)). \end{aligned}$$

This is exactly (5.20). ■

A.7 Computation of Gâteaux derivative (5.3)

In this section we prove (5.3), which states that

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

For ease of exposition, we 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 helpful¹²

$$\tau = \tilde{F}_\lambda \tilde{F}_\lambda^{-1}. \quad (\text{A.16})$$

To further simplify notation, write $q(\lambda) := \tilde{F}_\lambda^{-1}$. Then (A.16) 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

$$q'(0) = -\frac{-\tilde{F}(q(0)) + F(q(0))}{\tilde{f}(q(0))}. \quad (\text{A.17})$$

Notice that

$$\tilde{F}_\lambda|_{\lambda=0} = \tilde{F} \implies q(\lambda)|_{\lambda=0} = q(0) = \tilde{F}^{-1}. \quad (\text{A.18})$$

Substitute (A.18) into (A.17) to obtain

$$q'(0) = -\frac{-\tilde{F}(\tilde{F}^{-1}) + F(\tilde{F}^{-1})}{\tilde{f}(\tilde{F}^{-1})} = \frac{\tau - F(\tilde{F}^{-1})}{\tilde{f}(\tilde{F}^{-1})}. \quad (\text{A.19})$$

Notice that $q'(0)$ is exactly equal to the Gâteaux derivative from the definition in (5.2), since

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

¹²This “equality” may actually only be an inequality for some τ , but this is immaterial to the argument.

B Estimating the risk-neutral quantile function

B.1 Data description

To estimate the risk-neutral quantile curve for each point in time, I use daily option data from OptionMetrics covering the period 01-01-1996 until 12-31-2019. This consists of European Put and Call option data with time to expiration less than 500 days on the S&P 500 index. The option contract further contains data on the highest closing bid and lowest closing ask price and price of the forward contract on the underlying security. In addition, I obtain data on the daily risk-free rate from Kenneth French' website.¹³ Finally, stock price data on the closing price of the S&P 500 are obtained via WRDS.

Prior to estimating the martingale measure, I use an additional data cleaning procedure for the option data. All observations are dropped for which the highest closing bid price equals zero, as well as all option prices that violate no-arbitrage bounds. This is similar to the cleaning procedure of [Martin \(2017\)](#) and leaves a total of 16,624,104 option-day observations.

B.2 Estimating the risk-neutral quantile function

There is a substantial literature on how to extract the martingale measure from option prices. I follow [Figlewski \(2008\)](#), with some minor modifications emphasized below to estimate the conditional and unconditional risk-neutral quantile curve for tenors 30, 60, 90, 180 and 360 days.

- (i) Construct option prices: Use the midpoint of the highest closing bid and lowest closing ask price to obtain the option price.
- (ii) Convert option prices to Black-Scholes implied volatilities (IVs): Use out-of-the money put and call option prices to construct IVs, since these tend to be more liquid than in-the-money options. I regard a put option out-of-the-money if the strike price is less than the forward price on the underlying security. This is consistent with [Martin \(2017\)](#). Forward prices are provided by OptionMetrics.

¹³See http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html#Research

- (iii) Interpolate the IVs using a smoothing cubic spline: I use a smoothing cubic spline (Wahba, 1990) with 4 knots to interpolate the IVs at a dense set of equidistant strike prices

$$K \in \{K_{\min}, K_{\min} + \Delta, K_{\min} + 2\Delta, \dots, K_{\max}\} \quad \Delta = \frac{K_{\max} - K_{\min}}{500},$$

where K_{\min}, K_{\max} are the respective minimum and maximum strike prices observed in the sample. This differs from Figlewski (2008), who recommends a 4th degree smoothing spline and a single knot. The 4th degree spline in Figlewski (2008) comes from the necessity to obtain a smooth density function, which corresponds to the second derivative of the put-call-option price curve. Since I only need to estimate a CDF, a 3rd degree polynomial suffices and prevents overfitting. The use of 4 knots is arbitrary, as is the single knot in Figlewski (2008), but renders acceptable estimates of the martingale measure in most cases.

- (iv) Smooth the IV curve for at-the-money options: There tends to be a discontinuity in the smoothed IV curve for at-the-money options, since puts and calls at the same strike trade on slightly different IVs. Let F_{t+1} be the price of the forward contract, I consider all strike prices $K \in [K_{\text{low}}, K_{\text{high}}]$, where K_{low} is the lowest traded strike such that $F_{t+1}(1 - 0.02) \leq K_{\text{low}}$ and K_{high} the highest traded strike which satisfies $K_{\text{high}} \leq (1 + 0.02)F_{t+1}$. Following Figlewski (2008), I use a weighted average of $IV_{\text{put}}(K), IV_{\text{call}}(K)$ to estimate

$$IV(K) = \alpha IV_{\text{put}} + (1 - \alpha) IV_{\text{call}} \quad K \in [K_{\text{low}}, K_{\text{high}}],$$

where

$$\alpha = \frac{K_{\text{high}} - K}{K_{\text{high}} - K_{\text{low}}}.$$

- (v) Obtain the central part of the risk-neutral CDF: The smoothed Black-Scholes IVs are converted back to option prices. By the (nonparametric) result of Breeden and Litzenberger (1978)

$$1 + R_{f,t+1} \frac{\partial}{\partial K} \text{Call}_{t+1}(K) = \tilde{\mathbb{P}}_t(S_{t+1} \leq K) = \tilde{\mathbb{P}}_t \left(R_{m,t+1} \leq \frac{K}{S_t} \right). \quad (\text{B.1})$$

I convert put prices to call option prices using put-call parity. Let K_{n-1}, K_n, K_{n+1} be consecutive strike prices on the strike grid, then the partial derivative in (B.1) is approximated by

$$1 + R_{f,t+1} \left[\frac{C_{n+1} - C_{n-1}}{K_{n+1} - K_{n-1}} \right] \approx \tilde{\mathbb{P}}_t \left(R_{m,t+1} \leq \frac{K}{S_t} \right).$$

- (vi) Obtain the risk-neutral CDF in the tails: It remains to estimate the risk-neutral CDF for $K \leq K_2$ and $K \geq K_{499}$, which concerns the left and right tail of the distribution. I follow Figlewski (2008) and fit a generalized extreme value (GEV) distribution in the left and right tail. The GEV distribution function is given by (see De Haan and Ferreira (2007, p. 6)):

$$F_{\text{GEV}}(x) = \exp(-(1 + \gamma x)^{-1/\gamma}).$$

A location parameter μ and scale parameter σ can be introduced via the transformation

$$x = \frac{S_{t+1} - \mu}{\sigma}.$$

The function $F_{\text{GEV}}(S_{t+1}; \mu, \sigma, \gamma)$ effectively contains three parameters which will be calibrated to the implied risk-neutral CDF. For the left tail, I use $\tau = 0.02, \tau_2 = 0.03$ and $\tau_3 = 0.05$ and find the associated strikes from the empirical risk-neutral CDF, denoted by $K(\tau_1), K(\tau_2), K(\tau_3)$ respectively. Thereafter, I solve the following non-linear system of equations to match the quantiles of F_{GEV} with the empirical quantiles:

$$F_{\text{GEV}}(K(\tau_1)) = \tau_1$$

$$F_{\text{GEV}}(K(\tau_2)) = \tau_2$$

$$F_{\text{GEV}}(K(\tau_3)) = \tau_3.$$

I follow the same procedure to extend the risk-neutral CDF to the right tail, by matching the GEV function at the empirical quantile $\tau_1 = 0.95, \tau_2 = 0.97, \tau_3 = 0.98$. I discard tenors for which the 0.02 or 0.98-quantiles are not available, to prevent extrapolation error in the tail. I depart from Figlewski (2008) by matching only quantiles of the CDF, as opposed to

matching the shape of the PDF.

- (vii) Interpolate the risk-neutral CDFs for specific tenor: Typically, options with a specific tenor we want to estimate (e.g. 30 days) are not traded. To overcome this issue, I linearly interpolate the conditional risk-neutral curves with tenors closest to the tenor of interest. For example, suppose on a given day only options with tenor 20 and 45 days are traded and we are interested to obtain the 30-day risk-neutral CDF. In that case, the 20 and 45-day risk-neutrals CDFs are linearly interpolated so as to obtain the 30-day risk-neutral CDF. This is similar to [Martin \(2017, Appendix A\)](#). The conditional risk-neutral quantile curve is constructed as the left continuous inverse of the estimated risk-neutral CDF

$$\tilde{Q}_{t,\tau} = \inf \left\{ x \in \mathbb{R} : \tau \leq \tilde{\mathbb{P}}_t(R_{m,t+1} \leq x) \right\}.$$

- (viii) Estimate the unconditional risk-neutral CDF: Once the conditional risk-neutral measure has been estimated using the previous steps, I estimate the unconditional CDF via

$$\tilde{\mathbb{P}}_T(R_m \leq x) := \frac{1}{T} \sum_{t=1}^T \tilde{\mathbb{P}}_t(R_{m,t+1} \leq x).$$

- (ix) Estimate the unconditional risk-neutral quantile function: Given the unconditional risk-neutral CDF estimate in Step (viii), I estimate the unconditional risk-neutral quantile function with

$$\tilde{Q}_T(\tau) := \inf \left\{ x \in \mathbb{R} : \tau \leq \tilde{\mathbb{P}}_T(x) \right\}. \quad (\text{B.2})$$

C Detailed representative agent models

In this Section I show two results about representative agent models which are used in the paper. The first Section describes how to obtain the risk-neutral and physical CDF in the disaster risk model. Section [C.2](#) shows that the subjective crash risk probability derived by [Martin \(2017\)](#) under log preferences is identical to the crash probability I obtain building on the work of [Chabi-Yo and Loudis](#)

(2020).

C.1 Disaster risk probabilities

This section contains the details about the disaster risk figures shown in Figure 2. Let $F(\cdot)$ denote the CDF of Δc . Backus et al. (2011, p. 1976) show that

$$F(-b) = \sum_{j=0}^{\infty} \Phi\left(\frac{-b - \mu - j\theta}{\sqrt{\sigma^2 + j\nu^2}}\right) \cdot \frac{e^{-\kappa}\kappa^j}{j!}, \quad (\text{C.1})$$

where $\Phi(\cdot)$ is the CDF of a standard normal variable. From here it is straightforward to compute the physical CDF of return on equity, $R = \exp(\lambda\Delta c)$, by a change of variables and truncating the sum in (C.1). In practice I use 31 terms, which is very accurate. Secondly, to obtain the risk-neutral CDF, I use the result of Backus et al. (2011, p. 1987), that the risk-neutral distribution of Δc is the same as in (C.1), with new parameters

$$\tilde{\kappa} = \kappa e^{-\gamma\theta + (\gamma\nu)^2/2}, \quad \tilde{\theta} = \theta - \gamma\nu^2.$$

It follows that $\tilde{\kappa} > \kappa$ if $\theta < 0$ (more jumps) and $\tilde{\theta} < \theta$ (outcomes of jump is more negative on average). This explains the fat left tail of the risk-neutral distribution in Figure 2. The quantile bound, HJ bound and SDF volatility can now easily be calculated from the expression of the physical and risk-neutral distribution.

C.2 Crash probability with known utility

Chabi-Yo and Loudis (2020) 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) shows that

$$\mathbb{P}_t(R_{m,t+1} < \alpha) = \alpha \left[\text{Put}'_t(\alpha S_t) - \frac{\text{Put}_t(\alpha S_t)}{\alpha S_t} \right], \quad (\text{C.2})$$

where Put'_t is the derivative of the put option price curve seen as a function of the strike. Under log preferences and using (5.5), it follows that

$$\begin{aligned}
\mathbb{P}_t(R_{m,t+1} < \tilde{Q}_{t,\tau}) &= \tau + \frac{1}{R_{f,t+1}} \widetilde{\text{COV}}_t \left[\mathbb{1} \left(R_{m,t+1} \leq \tilde{Q}_{t,\tau} \right), R_{m,t+1} \right] \\
&= \tau + \frac{1}{R_{f,t+1}} \left(\tilde{\mathbb{E}}_t \left[\mathbb{1} \left(R_{m,t+1} \leq \tilde{Q}_{t,\tau} \right) R_{m,t+1} \right] - \tilde{\mathbb{E}}_t(R_{m,t+1}) \tilde{\mathbb{E}}_t \left(\mathbb{1} \left(R_{m,t+1} \leq \tilde{Q}_{t,\tau} \right) \right) \right) \\
&= \frac{1}{R_{f,t+1}} \tilde{\mathbb{E}}_t \left[\mathbb{1} \left(R_{m,t+1} \leq \tilde{Q}_{t,\tau} \right) R_{m,t+1} \right]. \tag{C.3}
\end{aligned}$$

The result now follows upon substituting $\tilde{Q}_\tau = \alpha$, since [Martin \(2017\)](#) shows that (C.3) equals the right hand side of (C.2).

D Other SDF bounds

The principal method to use quantiles to derive bounds on the volatility of the SDF can be applied to other well known bounds in the literature. In this Section I revisit some of these SDF bounds and show how the quantile relation can be used to obtain results akin to the quantile version of the HJ bound in [Theorem 2.1](#). For all the results to follow it is well known under which conditions the bounds are tight. For example, the log bound in [Section D.1](#) is known to bind for the growth-optimal portfolio. Under some conditions the growth-optimal portfolio is equal to the market portfolio. Using the quantile relation to bound the log of SDF could therefore refute the presumption that the market portfolio is growth optimal, if the quantile bound is significantly stronger. For convenience, recall the relation derived in the proof of [Theorem 2.1](#), which is used repeatedly in this Section to analyze other SDF bounds

$$\tau = R_{f,t+1} \mathbb{E}_t \left[M_{t+1} \mathbb{1} \left(R_{t+1} \leq \tilde{Q}_{t,\tau} \right) \right]. \tag{D.1}$$

D.1 Bound of [Bansal and Lehmann \(1997\)](#)

Here I consider a bound on the logarithm of the SDF. By an application of Jensen's inequality, we get

$$\begin{aligned} 0 = \log(1) &= \log \mathbb{E}_t [M_{t+1} R_{t+1}] \geq \mathbb{E}_t [\log M_{t+1}] + \mathbb{E}_t [\log R_{t+1}] \\ &\implies -\mathbb{E}_t [\log M_{t+1}] \geq \mathbb{E}_t [\log R_{t+1}]. \end{aligned}$$

This bound, together with its asset pricing implications, is analyzed in detail by [Bansal and Lehmann \(1997\)](#). It is known to bind for the market portfolio in a representative agent model with log utility. Applying a log transformation to [\(D.1\)](#), we obtain for any $\tau \in (0, 1)$

$$\log(\tau) = \log(R_{f,t+1}) + \log \left(\mathbb{E}_t \left[M_{t+1} \mathbb{1} \left(R_{t+1} \leq \tilde{Q}_{t,\tau} \right) \right] \right).$$

Use Jensen's inequality in a similar vein as above and rearranging gives

$$-\mathbb{E}_t [\log(M_{t+1})] \geq \log(R_{f,t+1}) + \mathbb{E}_t \left[\log \left(\mathbb{1} \left(R_{t+1} \leq \tilde{Q}_{t,\tau} \right) \right) \right] - \log(\tau).$$

Taking expectations on both sides also renders an unconditional version.

D.2 Bound of [Snow \(1991\)](#)

[Snow \(1991\)](#) derives a continuum of bounds of higher order moments on the SDF. In somewhat simplified form, the idea is to use Hölder's inequality to the defining SDF equation

$$1 = \mathbb{E}_t [M_{t+1} R_{t+1}] \leq \mathbb{E}_t [M_{t+1}^p]^{\frac{1}{p}} \mathbb{E}_t [R_{t+1}^q]^{\frac{1}{q}},$$

for Hölder exponents $\frac{1}{p} + \frac{1}{q} = 1$ and $p > 1$. Rearranging gives the restriction on the p -th norm of the SDF

$$\mathbb{E}_t [M_{t+1}^p]^{\frac{1}{p}} \geq \mathbb{E}_t [R_{t+1}^q]^{-\frac{1}{q}}.$$

The quantile relation (D.1) can similarly be exploited by applying Hölder's inequality on the right hand side. This gives

$$\mathbb{E}_t [M_{t+1}^p]^{\frac{1}{p}} \geq \left(\frac{\tau}{R_{f,t+1}} \right) \mathbb{E}_t \left[\mathbb{1} \left(R_{t+1} \leq \tilde{Q}_{t,\tau} \right) \right]^{-\frac{1}{q}}.$$

D.3 Bound of Liu (2020)

Liu (2020) develops a continuum of bounds which are based on different moments of the SDF. In particular

$$\mathbb{E}_t [M_{t+1}^s] \begin{cases} \leq \mathbb{E}_t \left[R_{t+1}^{-\frac{s}{1-s}} \right]^{1-s}, & \text{if } s \in (0, 1). \\ \geq \mathbb{E}_t \left[R_{t+1}^{-\frac{s}{1-s}} \right]^{1-s}, & \text{if } s \in (-\infty, 0). \end{cases} \quad (\text{D.2})$$

The proof, as in Liu (2020), follows from an application of the *reverse Hölder inequality*.¹⁴ Equality occurs for the return which satisfies

$$\log M_{t+1} = -\frac{1}{1-s} \log R_{t+1} + \text{Constant}.$$

The quantile relation can only be used to obtain the upper bound part in (D.2), since the reverse Hölder inequality requires almost sure positivity of $\mathbb{1}(R_{t+1} \leq \tilde{Q}_{t,\tau})$ to prove the lower bound. For $p \in (1, \infty)$, apply the reverse Hölder inequality to the relation (D.1) to obtain

$$\tau = R_{f,t+1} \mathbb{E}_t \left[M_{t+1} \mathbb{1} \left(R_{t+1} \leq \tilde{Q}_{t,\tau} \right) \right] \geq R_{f,t+1} \mathbb{E} \left[M_{t+1}^{\frac{-1}{p-1}} \right]^{1-p} \mathbb{E}_t \left[\mathbb{1} \left(R_{t+1} \leq \tilde{Q}_{t,\tau} \right)^{\frac{1}{p}} \right]^p$$

Rearranging and using $s := -\frac{1}{p-1} \in (-\infty, 0)$ yields

$$\mathbb{E}_t [M_{t+1}^s] \geq \left(\frac{\tau}{R_{f,t+1}} \right)^s \mathbb{E}_t \left[\mathbb{1} \left(R_{t+1} \leq \tilde{Q}_{t,\tau} \right) \right]^{1-s}.$$

¹⁴The reverse Hölder inequality states that for any $p \in (1, \infty)$ and measure space (S, Σ, μ) that satisfies $\mu(S) > 0$. Then for all measurable real- or complex-valued functions f and g on S such that $g(s) \neq 0$ for μ -almost all $s \in S$, $\|fg\|_1 \geq \|f\|_{\frac{1}{p}} \|g\|_{\frac{-1}{p-1}}$.

E Von Mises approximation: evidence from Black-Scholes

This section illustrates the accuracy of the quantile approximation in (5.10) in a discretized version of the Black-Scholes model with changing parameters. Specifically, I assume the following model

$$R_{t+1} = \exp\left(\left(\mu_t - \frac{1}{2}\sigma_t^2\right)\lambda + \sigma_t\lambda Z_{t+1}\right) \quad (\text{E.1})$$

$$Z_{t+1} \sim N(0, 1)$$

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

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

The returns under risk-neutral dynamics are given by

$$R_{t+1} = \exp\left(\left(r_t - \frac{1}{2}\sigma_t^2\right)\lambda + \sigma_t\lambda Z_{t+1}\right) \quad (\text{E.2})$$

$$r_t \sim \text{UNIF}[0, 0.03]. \quad (\text{E.3})$$

Finally, assume that all parameters are IID over time and that options are priced according to the Black-Scholes formula, conditional on time t . In this setup, it is fruitless to use historical data to predict future quantiles, since parameters change unpredictably over time. We use $\lambda = 1/12$ to mimic the monthly application in this paper. It is assumed that the risk-neutral quantile function is known at the start of period t , as it is in the real world, using the result of Breeden and Litzenberger (1978). I use the risk-neutral quantile function to calculate $LRB_t(\tau)$ at time t . Then, following the approximation in (5.10), the physical quantile function is estimated by

$$\hat{Q}_{t,\tau} = \tilde{Q}_{t,\tau} + \frac{LRB_t(\tau)}{f_t(\tilde{Q}_{t,\tau})}. \quad (\text{E.4})$$

We take 3,000 return observations, which are generated according to (E.1). This exercise is repeated 1,000 times. To assess the accuracy of the approximation in (E.4), I use several metrics. For every sample, I estimate a quantile regression

Table 6: Simulation results

	$\mathbb{E}\widehat{\beta}_0(\tau)$	$\mathbb{E}\widehat{\beta}_1(\tau)$	$Q > \widehat{Q}$	$\rho(Q, \widehat{Q})$	$H_0 : \widehat{\beta}_0(\tau) = 0$	$H_0 : \widehat{\beta}_1(\tau) = 1$	$H_0 : [\widehat{\beta}_0(\tau), \widehat{\beta}_1(\tau)] = [0, 1]$
$\tau = 0.01$	0.01	0.99	0.85	1	0.94	0.96	0.8
$\tau = 0.05$	-0.03	1.04	0.69	0.99	0.9	0.89	0.66
$\tau = 0.1$	-0.06	1.07	0.64	0.99	0.78	0.76	0.47

Note: $\mathbb{E}\widehat{\beta}_0(\tau)$ denotes the average quantile regression estimate of $\widehat{\beta}_0(\tau)$ and likewise $\mathbb{E}\widehat{\beta}_1(\tau)$ shows it for $\widehat{\beta}_1(\tau)$. $Q > \widehat{Q}$ shows the fraction of times the true physical quantile is larger than our predicted quantile. Columns $H_0 : \widehat{\beta}_0(\tau) = 0$ and $H_0 : \widehat{\beta}_1(\tau) = 1$ report the fraction of times the individual null hypotheses $\beta_0(\tau) = 0, \beta_1(\tau) = 1$ are not rejected. The last column reports the fraction of times the joint null hypothesis is not rejected.

of the form

$$Q_\tau(R_{t+1}) = \beta_0(\tau) + \beta_1(\tau)\widehat{Q}_{t,\tau},$$

where $\widehat{Q}_{t,\tau}$ comes from (E.4). The first two columns in Table 6 report the average values of the quantile regression estimates across the 1,000 simulations. The means are rather close to 0 and 1 respectively for all quantiles. If (5.4) is a good approximation, one expects $Q_{t,\tau} > \widehat{Q}_{t,\tau}$, since $LRB_t(\tau) \leq \tau - \mathbb{P}_t(\widehat{Q}_{t,\tau})$. The third column in Table 6 shows this happens for the majority of samples. The fourth column shows the correlation between $Q_{t,\tau}$ and $\widehat{Q}_{t,\tau}$, which is very close to one, and corroborates the view that the approximation is quite accurate. Columns four and five document the percentage of non rejection of H_0 , which is indeed quite high. The last column considers non rejection of the joint null hypothesis, which is also high except for the $\tau = 0.1$ quantile. Overall, Table 6 suggests that (E.4) is a highly accurate predictor of the physical quantile function.

Example E.1. Let us illustrate the von Mises approximation (5.4) in the Black-Scholes model with fixed parameters: $\lambda = 1$ (one year), $\mu = 0.08, r = 0.02, \sigma = 0.2$.¹⁵ We can explicitly calculate $F^{-1}, \widetilde{F}^{-1}$ and \widetilde{f} owing to the lognormal assumption. Figure 6 shows the risk-neutral quantile function (in green), von Mises approximation (5.4) (in blue) and the true physical quantile function (in red). Observe that the approximation (5.4) is very accurate in this case.

E.1 Bias in quantile regression

In the empirical application, we have to estimate $\widehat{Q}_{t,\tau}, \widetilde{f}(\cdot)$ and $LRB_t(\tau)$. Therefore, the estimated coefficients in the quantile regression are biased, due to mea-

¹⁵For illustrative purposes, I use $\lambda = 1$, instead of $\lambda = 1/12$, otherwise the physical quantile function and von Mises approximation are indistinguishable.

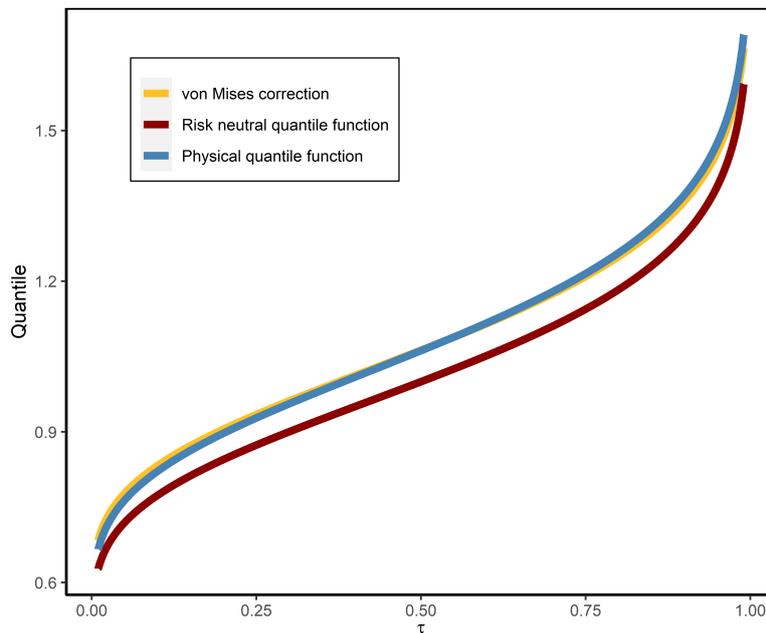


Figure 6: Illustration of quantile approximation (E.4) in the Black-Scholes model.

surement error in the covariate. I present simulation evidence which shows that the bias is small in finite samples.

The setup is as follows. We simulate returns according to model (E.1) and assume that options are priced with Black and Scholes (1973) at the start of period t . We want to calculate the von Mises approximation for a maturity of 90 days. As in the 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.¹⁶ These numbers are roughly consistent with the latter part of our empirical sample. The procedure is repeated for a total of 1,000 time periods. For the entire sample, I compare

¹⁶So on average there will be 500 options with maturity 85 days and 500 with maturity 97 days.

the estimated and analytical von Mises term, which respectively are given by

$$VM_{t,\tau}^e := \widehat{Q}_{t,\tau} + \frac{\widehat{LRB}_t(\tau)}{\widehat{f}_t(\widehat{Q}_{t,\tau})}$$

$$VM_{t,\tau}^a := \widetilde{Q}_{t,\tau} + \frac{LRB_t(\tau)}{\widetilde{f}_t(\widetilde{Q}_{t,\tau})}.$$

The hats in VM^e signify that the risk-neutral quantile, pdf and lower bound are estimated from the available options at time t , using the procedure in Appendix B.2. The terms in VM^a are obtained from the known analytical expression of the risk-neutral distribution (recall (E.2)). I then use quantile regression to estimate the models

$$Q(R_{t+1}) = \widehat{\beta}_0(\tau) + \widehat{\beta}_{1,e}(\tau)VM_{t,\tau}^e$$

$$Q(R_{t+1}) = \widehat{\beta}_0(\tau) + \widehat{\beta}_{1,a}(\tau)VM_{t,\tau}^a.$$

I use the ratio $\widehat{\beta}_{1,e}/\widehat{\beta}_{1,a}$ to measure the relative bias in the sample. This is repeated 500 times to get a distribution of the relative bias. Figure 7 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 von Mises term is limited.

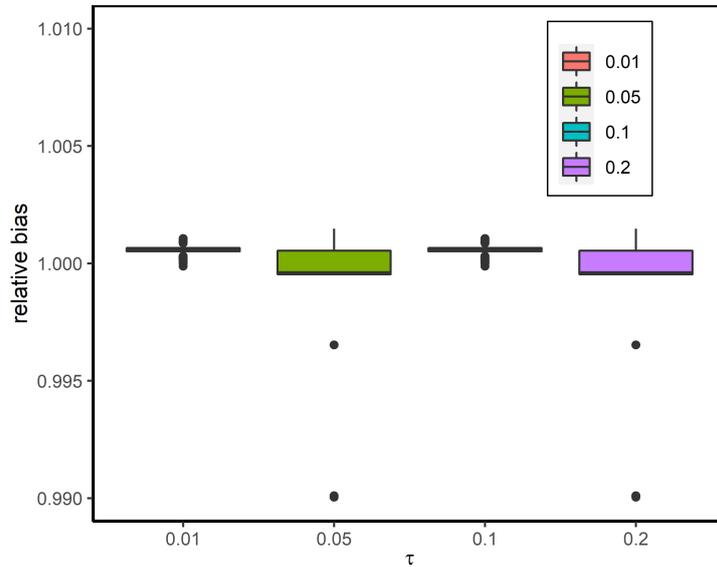


Figure 7: Bias in quantile regression resulting from measurement error.

References

- Ait-Sahalia, Y., Fan, J., and Li, Y. (2013). The leverage effect puzzle: Disentangling sources of bias at high frequency. *Journal of Financial Economics*, 109(1):224–249.
- Ait-Sahalia, Y. and Lo, A. W. (1998). Nonparametric estimation of state-price densities implicit in financial asset prices. *The Journal of Finance*, 53(2):499–547.
- Ait-Sahalia, Y. and Lo, A. W. (2000). Nonparametric risk management and implied risk aversion. *Journal of econometrics*, 94(1-2):9–51.
- Almeida, C. and Garcia, R. (2012). Assessing misspecified asset pricing models with empirical likelihood estimators. *Journal of Econometrics*, 170(2):519–537.
- Alvarez, F. and Jermann, U. J. (2005). Using asset prices to measure the persistence of the marginal utility of wealth. *Econometrica*, 73(6):1977–2016.
- Backus, D., Chernov, M., and Martin, I. (2011). Disasters implied by equity index options. *The Journal of Finance*, 66(6):1969–2012.

- Backus, D., Chernov, M., and Zin, S. (2014). Sources of entropy in representative agent models. *The Journal of Finance*, 69(1):51–99.
- Bai, J. (2003). Inferential theory for factor models of large dimensions. *Econometrica*, 71(1):135–171.
- Bakshi, G., Madan, D., and Panayotov, G. (2010). Returns of claims on the upside and the viability of U-shaped pricing kernels. *Journal of Financial Economics*, 97(1):130–154.
- Bansal, R., Kiku, D., and Yaron, A. (2012). An empirical evaluation of the long-run risks model for asset prices. *Critical Finance Review*, (1):183–221.
- Bansal, R. and Lehmann, B. N. (1997). Growth-optimal portfolio restrictions on asset pricing models. *Macroeconomic dynamics*, 1(2):333–354.
- Bansal, R. and Yaron, A. (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *The journal of Finance*, 59(4):1481–1509.
- Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, 121(3):823–866.
- Bates, D. S. (1991). The crash of 87: was it expected? The evidence from options markets. *The Journal of Finance*, 46(3):1009–1044.
- Bates, D. S. (2000). Post-'87 crash fears in the s&p 500 futures option market. *Journal of econometrics*, 94(1-2):181–238.
- Bates, D. S. (2008). The market for crash risk. *Journal of Economic Dynamics and Control*, 32(7):2291–2321.
- Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654.
- Bollerslev, T. and Todorov, V. (2011). Tails, fears, and risk premia. *The Journal of Finance*, 66(6):2165–2211.
- Borovička, J., Hansen, L. P., and Scheinkman, J. A. (2016). Misspecified recovery. *The Journal of Finance*, 71(6):2493–2544.

- Breedon, D. T. and Litzenberger, R. H. (1978). Prices of state-contingent claims implicit in option prices. *Journal of business*, pages 621–651.
- Broadie, M., Chernov, M., and Johannes, M. (2009). Understanding index option returns. *The Review of Financial Studies*, 22(11):4493–4529.
- Campbell, J. Y. and Thompson, S. B. (2008). Predicting excess stock returns out of sample: Can anything beat the historical average? *The Review of Financial Studies*, 21(4):1509–1531.
- Chabi-Yo, F. and Loudis, J. (2020). The conditional expected market return. *Journal of Financial Economics*.
- Cochrane, J. H. (2005). *Asset pricing: Revised edition*. Princeton university press.
- Coval, J. D. and Shumway, T. (2001). Expected option returns. *The Journal of Finance*, 56(3):983–1009.
- Danielsson, J. and De Vries, C. G. (2000). Value-at-risk and extreme returns. *Annales d’Economie et de Statistique*, pages 239–270.
- De Haan, L. and Ferreira, A. (2007). *Extreme value theory: an introduction*. Springer Science & Business Media.
- Engle, R. F. and Manganelli, S. (2004). Caviar: Conditional autoregressive value at risk by regression quantiles. *Journal of Business & Economic Statistics*, 22(4):367–381.
- Epstein, L. G. and Zin, S. E. (1989). Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, (57):937–969.
- Figlewski, S. (2008). Estimating the implied risk neutral density.
- Gabaix, X. (2012). Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *The Quarterly journal of economics*, 127(2):645–700.
- Gouriéroux, C. and Jasiak, J. (2008). Dynamic quantile models. *Journal of Econometrics*, 147(1):198–205.

- Goyal, A. and Welch, I. (2008). A comprehensive look at the empirical performance of equity premium prediction. *The Review of Financial Studies*, 21(4):1455–1508.
- Gregory, K. B., Lahiri, S. N., and Nordman, D. J. (2018). A smooth block bootstrap for quantile regression with time series. *The Annals of Statistics*, 46(3):1138–1166.
- Hansen, L. P. and Hodrick, R. J. (1980). Forward exchange rates as optimal predictors of future spot rates: An econometric analysis. *Journal of political economy*, 88(5):829–853.
- Hansen, L. P. and Jagannathan, R. (1991). Implications of security market data for models of dynamic economies. *Journal of political economy*, 99(2):225–262.
- Hodrick, R. J. (1992). Dividend yields and expected stock returns: Alternative procedures for inference and measurement. *The Review of Financial Studies*, 5(3):357–386.
- Hsieh, F., Turnbull, B. W., et al. (1996). Nonparametric and semiparametric estimation of the receiver operating characteristic curve. *Annals of statistics*, 24(1):25–40.
- Jackwerth, J. C. (2000). Recovering risk aversion from option prices and realized returns. *The Review of Financial Studies*, 13(2):433–451.
- Jackwerth, J. C. and Menner, M. (2020). Does the ross recovery theorem work empirically? *Journal of Financial Economics*, 137(3):723–739.
- Koenker, R. and Bassett, G. (1978). Regression quantiles. *Econometrica: journal of the Econometric Society*, pages 33–50.
- Koenker, R. and Machado, J. A. (1999). Goodness of fit and related inference processes for quantile regression. *Journal of the american statistical association*, 94(448):1296–1310.
- Liu, Y. (2020). Index option returns and generalized entropy bounds. *Journal of Financial Economics*.

- Martin, I. (2017). What is the expected return on the market? *The Quarterly Journal of Economics*, 132(1):367–433.
- Martin, I. and Gao, C. (2021). Volatility, valuation ratios, and bubbles: An empirical measure of market sentiment. *Journal of Finance*, forthcoming.
- Qin, L., Linetsky, V., and Nie, Y. (2018). Long forward probabilities, recovery, and the term structure of bond risk premiums. *The Review of Financial Studies*, 31(12):4863–4883.
- Rietz, T. A. (1988). The equity risk premium a solution. *Journal of monetary Economics*, 22(1):117–131.
- Rosenberg, J. V. and Engle, R. F. (2002). Empirical pricing kernels. *Journal of Financial Economics*, 64(3):341–372.
- Ross, S. A. (1976). Options and efficiency. *The Quarterly Journal of Economics*, 90(1):75–89.
- Ross, S. A. (2015). The recovery theorem. *The Journal of Finance*, 70(2):615–648.
- Serfling, R. J. (2009). *Approximation theorems of mathematical statistics*, volume 162. John Wiley & Sons.
- Snow, K. N. (1991). Diagnosing asset pricing models using the distribution of asset returns. *The Journal of Finance*, 46(3):955–983.
- Stutzer, M. (1995). A bayesian approach to diagnosis of asset pricing models. *Journal of Econometrics*, 68(2):367–397.
- Van der Vaart, A. W. (2000). *Asymptotic statistics*, volume 3. Cambridge university press.
- von Mises, R. (1947). On the asymptotic distribution of differentiable statistical functions. *The annals of mathematical statistics*, 18(3):309–348.
- Wachter, J. A. (2013). Can time-varying risk of rare disasters explain aggregate stock market volatility? *The Journal of Finance*, 68(3):987–1035.
- Wahba, G. (1990). *Spline models for observational data*. SIAM.