JOURNAL OF APPLIED ECONOMETRICS
J. Appl. Econ. **32**: 192–217 (2017)
Published online 1 March 2016 in Wiley Online Library
(wileyonlinelibrary.com) DOI: 10.1002/jae.2515

STATE PRICES OF CONDITIONAL QUANTILES: NEW EVIDENCE ON TIME VARIATION IN THE PRICING KERNEL

Konstantinos Metaxoglou[a] AND Aaron Smith[b]*

[a] *Department of Economics Carleton University Ottawa Ontario Canada*
[b] *Department of Agricultural and Resource Economics University of California Davis CA USA*

SUMMARY

We develop a set of statistics to represent the option-implied stochastic discount factor and we apply them to S&P 500 returns between 1990 and 2012. Our statistics, which we call state prices of conditional quantiles (SPOCQ), estimate the market's willingness to pay for insurance against outcomes in various quantiles of the return distribution. By estimating state prices at conditional quantiles, we separate variation in the shape of the pricing kernel from variation in the probability of a particular event. Thus, without imposing strong assumptions about the distribution of returns, we obtain a novel view of pricing-kernel dynamics. We document six features of SPOCQ for the S&P 500. Most notably, and in contrast to recent studies, we find that the price of downside risk decreases when volatility increases. Under a standard asset pricing model, this result implies that most changes in volatility stem from fluctuations in idiosyncratic risk. Consistent with this interpretation, no known systematic risk factors such as consumer sentiment, liquidity or macroeconomic risk can account for the negative relationship between the price of downside risk and volatility.

Received 17 November 2014; Revised 12 November 2015


Supporting information may be found in the online version of this article.

1. INTRODUCTION

In a present value model, stock price changes stem either from changes in expected future dividends or changes in the discount rate. If the discount rate is constant, then price changes must reflect changes in expected future dividends. Early work in finance showed that stock prices vary much more than can be explained by rational expectations about future dividends. Accordingly, discount rate variation has become a focal point in asset pricing research.

In this paper, we provide a new perspective on discount-rate variation by developing a set of statistics to represent the option-implied stochastic discount factor. To generate these statistics, we first use a time series model to estimate quantiles of the conditional distribution of returns. Then, taking these estimates to represent rational expectations about return quantiles, we use options prices to infer the implied discount factor in interquantile ranges of the return distribution. Specifically, we estimate the prices of securities that pay 1 dollar in the event that the asset return falls in, for example, the bottom 10% of the conditional distribution, and zero otherwise. We call our statistics state prices of conditional quantiles (SPOCQ) and apply them to S&P 500 monthly returns between 1990 and 2012.

We use quantiles for three reasons. First, we can obtain a parsimonious model of the important features of the conditional return distribution without imposing the assumptions required to fit the entire distribution. Second, quantiles often provide a more intuitive and useful characterization of tail behavior than moments. Value-at-risk (VaR) is one prominent example; a VaR statistic provides a

* Correspondence to: Aaron Smith, Department of Agricultural and Resource Economics, University of California, Davis, CA, USA E-mail: adsmith@ucdavis.edu

clearer picture of tail risk than a skewness or kurtosis statistic. Third, and perhaps most important, time series variation in SPOCQ reveals information about discount rate variation that is obscured in other statistics.

For example, consider the event that the 4-week return on the S&P 500 is less than −4.4%, which occurs in 10% of the months in our sample. The state price of this event can be expressed as the event probability multiplied by the discount factor traders apply to the event. The state price increases when volatility increases because high volatility makes a return less than −4.4% more likely, even if the stochastic discount factor remains unchanged. A return in the bottom *conditional* decile, on the other hand, always has a 10% probability of occurring. As a result, variation in the associated SPOCQ comes only from changes in the discount applied to that event. By holding constant the event probability, we gain insight into the changes in discount rates over time.

We document six features of discount rate variation for the S&P 500. First, left-tail returns have larger average SPOCQ than right-tail returns, which is consistent with standard notions of risk aversion; hence traders apply the largest discount to large negative returns.

Second, the SPOCQ for the top quartile typically exceeds its counterpart for the third quartile, implying that traders discount top-quartile returns more heavily than third-quartile returns. This finding is consistent with Jackwerth (2000) and, more recently, with Bakshi *et al.* (2010), who show that the pricing kernel is non-monotonic when projected on S&P 500 returns. This stylized fact of non-monotonicity has been labeled the pricing kernel puzzle—it contradicts standard asset pricing models under risk aversion, which predict that state prices for all lower-quantile events exceed those for higher-quantile events.

Third, SPOCQ exhibits substantial month-to-month variation. However, the monthly shocks in SPOCQ do not forecast returns or quantile hits, which suggests that the observed 'noise' is not attributable to estimation error in the quantiles or to changes in expected returns. Fourth, the left-tail SPOCQ peaked in months when major events occurred in the recent financial crisis, notably the fourth quarter of 2008 (Lehman Brothers bankruptcy), fourth quarter of 2007 (beginning of US credit crunch and recession), and two quarters in 2010–11 coinciding with the European sovereign debt crises. Hence these events coincided with a high estimated price of risk. Fifth, increases in the dividend yield and the term spread are associated with clockwise rotations in the pricing kernel, thereby increasing expected returns. This pattern is consistent with these variables capturing risk factors for equities. However, we find no such evidence for the default spread, CAY, consumer sentiment or the economic policy uncertainty index.

Finally, and most notably, we find that high volatility in stock returns is associated with lower state prices for bottom-quartile returns and higher state prices for top-quartile returns. Figure 1 illustrates this finding by plotting the bottom- and top-quartile SPOCQ statistics against the log of volatility—the two slopes are significantly different from each other at the 1% level. The negative correlation in the bottom quartile contradicts recent models that specify the pricing kernel as a linear function of squared returns (see Christoffersen *et al.* (2013). In those models, state prices in both tails of the return distribution increase with volatility.

Our results support the hypothesis that changes in volatility during our sample were predominantly *idiosyncratic* to the S&P 500 and thus unrelated to the pricing kernel. By definition, an increase in idiosyncratic volatility does not affect expected returns. Moreover, an increase in idiosyncratic volatility decreases the signal-to-noise ratio in S&P 500 returns. If returns contain relatively more noise then they convey less information about systematic risk. We show that, as a result, the left-tail SPOCQ decreases when idiosyncratic volatility increases. In contrast, an increase in systematic volatility increases the left-tail SPOCQ because it causes traders to apply a larger discount to left-tail returns.

Our results are consistent with the literature examining the return–volatility relationship. Beginning with French *et al.* (1987), numerous papers have found weak and inconsistent association between volatility and expected returns (see Yu and Yuan, 2011, for a summary). If increases in volatility tend to

DOI: 10.1002/jae

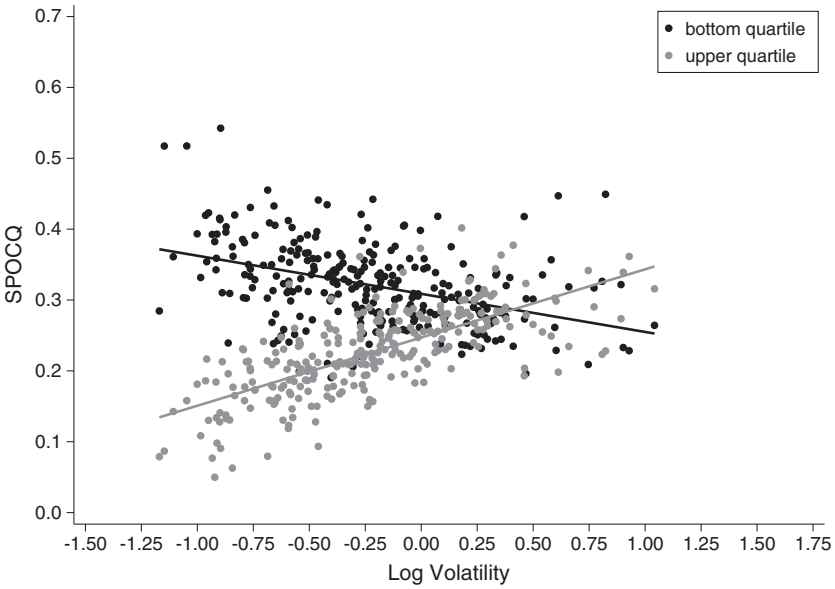

Figure 1. Scatter plot of the SPOCQ against the logarithm of volatility. We measure volatility using the square root of *realized* continuous variation of daily returns over the last 20 trading days. We compute this statistic using the method in Bollerslev and Todorov (2011), which we describe in online Appendix A.6

reflect greater *idiosyncratic* rather than *systematic* variation, then volatility should only weakly predict returns. Most of the papers in this literature cover time periods prior to our sample, which suggests that our results are not specific to the period since 1990.

The remainder of the paper is organized as follows. In Section 2 we provide a framework for SPOCQ and how to interpret its variation over time. We describe in Section 3 how we estimate the SPOCQ components, and we discuss our estimation results in Section 4. Section 5 describes the salient features of the SPOCQ time series, and in Section 6 we present regression results that relate SPOCQ to known risk factors. Section 7 concludes.

2. STATE PRICES OF CONDITIONAL QUANTILES

2.1. Conceptual Framework

In dynamic equilibrium models, the price of an asset equals the expected value of discounted future payoffs on the asset. Let S_t denote the price of the asset at time t, S_T the payoff on that asset at time $T > t$, and $M_{t,T}$ denote the stochastic discount factor between t and T. Suppose the state of the economy at t can be described by a vector W_t. In equilibrium, the asset price is given by

$$S_t = E_t \left[M_{t,T} S_T \right] \tag{1}$$

where $E_t[\cdot]$ denotes the expectation conditional on W_t.

At time t, the researcher observes S_t and the prices of any derivatives defined by payoffs on the asset, but not W_t. Thus we focus on the risk-neutral distribution implied by the observed asset and derivative prices, which is the risk-neutral distribution of returns on the asset after integrating out the unobserved component of the state space. Using the law of iterated expectations, the expression in equation (1) becomes

DOI: 10.1002/jae

$$1 = E_t\left[M_{t,T}R_{t,T}\right] = E_t\left[E_t\left[M_{t,T}|R_{t,T}\right]R_{t,T}\right] = \int E_t\left[M_{t,T}|R_{t,T}\right]R_{t,T}\mathrm{d}F_t\left(R_{t,T}\right) \quad (2)$$

where $R_{t,T} \equiv S_T/S_t$ and $E_t\left[\cdot|R_{t,T}\right]$ is the expectation conditional on $\{W_t, R_{t,T}\}$. Multiplying and dividing the last expression by $E_t\left[M_{t,T}\right]$ produces

$$1 = E_t\left[M_{t,T}\right]\int \frac{E_t\left[M_{t,T}|R_{t,T}\right]}{E_t\left[M_{t,T}\right]}R_{t,T}\mathrm{d}F_t(R_{t,T}) = E_t\left[M_{t,T}\right]\int R_{t,T}\mathrm{d}F_t^*(R_{t,T}) \quad (3)$$

Thus the risk-neutral conditional distribution of the asset return is given by

$$F_t^*(R) \equiv \int_{-\infty}^{R}\frac{E_t[M_{t,T}|R_{t,T}]}{E_t[M_{t,T}]}\mathrm{d}F_t(R_{t,T}) \equiv E_t\left[M_{t,T}^*|R_{t,T} \le R\right]F_t(R) \quad (4)$$

where we define $M_{t,T}^*$ as

$$M_{t,T}^* \equiv M_t^*(R_{t,T}) \equiv \frac{E_t\left[M_{t,T}|R_{t,T}\right]}{E_t\left[M_{t,T}\right]} \quad (5)$$

At time t, i.e. taking the current state of the world (W_t) as given, $M_t^*(R)$ is the discount applied to the return outcome R. The steps in equations (2)–(5) are similar to the projection of the pricing kernel onto the payoffs of a tradable asset in Engle and Rosenberg (2002). Therefore, $M_{t,T}^*$ can be labeled the projected pricing kernel.

Equation (4) decomposes the risk-neutral conditional distribution into two components and thereby reveals two sources of time variation. The first source of variation is changes in the future return distribution, $F_t(R)$. The second source of variation is changes in the price of risk, $M_t^*(R)$. We focus on the second source, which we isolate by evaluating $F_t^*(R)$ at conditional quantiles of the asset returns. Specifically, we define the conditional quantile $q_t(\theta_j)$ such that $F_t(q_t(\theta)) = \theta_j$. The state price of the event $R_{t,T} \le q_t(\theta_j)$, which occurs with fixed probability θ_j, is then given by

$$F_t^*\left(q_t(\theta_j)\right) = E_t\left[M_{t,T}^*|R_{t,T} \le q_t(\theta_j)\right]\theta_j \quad (6)$$

Equation (6) is an expression for a state price reflecting the market's willingness to pay for insurance against an event with a fixed probability.

We now define the state price of conditional quantiles (SPOCQ):

$$\mathrm{SPOCQ}_{t,T}\left(\theta_{j-1},\theta_j\right) = F_t^*\left(q_t(\theta_j)\right) - F_t^*\left(q_t(\theta_{j-1})\right) \quad (7)$$

Figure 2 illustrates how we obtain SPOCQ following a two-step approach. In the first step, we invert the physical distribution of returns to find the quantiles $q_t\left(\theta_{j-1}\right)$ and $q_t\left(\theta_j\right)$. In the second step, we evaluate the risk-neutral distribution at these quantiles.

Equivalently, using equation (6), we can write SPOCQ as

$$\mathrm{SPOCQ}_{t,T}\left(\theta_{j-1},\theta_j\right) = \int_{q_t(\theta_{j-1})}^{q_t(\theta_j)} M_{t,T}^*\mathrm{d}F_t(R_{t,T}) = \left(\theta_j - \theta_{j-1}\right)E_t\left[M_{t,T}^*|\Re_{t,T}^{j-1,j}\right] \quad (8)$$

where $\theta_j > \theta_{j-1}$ and $\Re_{t,T}^{j-1,j}$ denotes the states of the world at time T for which $q_t(\theta_{j-1}) \le R_{t,T} \le q_t(\theta_j)$. SPOCQ is the market's willingness to pay to receive a dollar in the event that the future return falls between the θ_{j-1} and θ_j quantiles. It equals the probability of this event multiplied by

DOI: 10.1002/jae

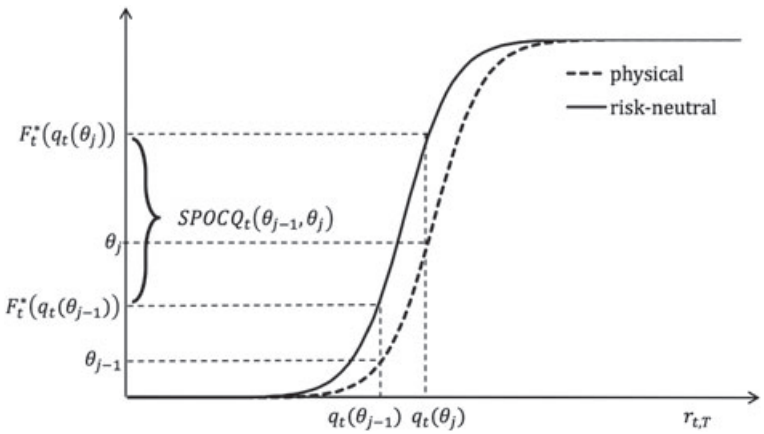

Figure 2. State price of conditional quantiles (SPOCQ) interpretation. We obtain SPOCQ at a particular quantile (θ_j) by first inverting the physical return distribution to find the quantile $q_t(\theta_j)$ and then evaluating the risk-neutral distribution at $q_t(\theta_j)$.

the average of the projected pricing kernel conditional on this event. The time variation in SPOCQ is driven entirely by the willingness to pay for insurance against this event because the probability of a return in this interval is fixed. Under risk neutrality, this state price would never change, and it would always equal the probability of the event occurring, $\theta_j - \theta_{j-1}$.

If we divide SPOCQ by the event probability, $\theta_j - \theta_{j-1}$, we obtain the mean of the projected pricing kernel conditional on the relevant return quantiles. We refer to $E_t[M^*_{t,T}|\Re^{j-1,j}_{t,T}]$ as the *quantile pricing kernel* (QPK) because it equals the mean of the pricing kernel conditional on the interquantile range of the return distribution dictated by θ_j and θ_{j-1}.

2.2. Interpreting Time Variation in SPOCQ

When projecting the pricing kernel onto returns for a broad index like the S&P 500, as we do in this paper, it is tempting to assume that the projection entails only a small loss in information. In the language of CAPM, it is tempting to assume that the S&P 500 is a good proxy for the market portfolio. In this section, we show how the time variation in SPOCQ can reveal the extent to which this assumption holds. We develop our argument using a single-factor model in which the stochastic discount factor is quadratic in a single state variable. This formulation captures the empirical regularity that the projected pricing kernel is non-monotonic, a finding known as the pricing kernel puzzle (Jackwerth, 2000).

Suppose the stochastic discount factor is given by

$$M_{t,T} = \frac{1}{R^f_{t,T}} - W_T + \gamma_t \left(W_T^2 - \sigma^2_{W,t}\right) \tag{9}$$

where $W_T|W_t \sim N\left(0, \sigma^2_{W,t}\right)$, $\gamma_t \geq 0$, and the risk-free rate $R^f_{t,T}$ may vary over time. This formulation imposes the equilibrium condition $E_t[M_{t,T}] = 1/R^f_{t,T}$. Stock returns include both a systematic component (W_T) and an idiosyncratic component ($\varepsilon_{t,T}$), and are generated by

$$R_{t,T} = \mu_t + \beta_t W_T + \varepsilon_{t,T} \tag{10}$$

where $\varepsilon_{t,T} \sim N\left(0, \sigma^2_{\varepsilon,t}\right)$ and $E_t[W_T \varepsilon_{t,T}] = 0$. We impose conditional normality on W_T and $\varepsilon_{t,T}$ in order to obtain closed-form expressions.

DOI: 10.1002/jae

In equilibrium, the equity risk premium is

$$E_t[R_{t,T}] - R^f_{t,T} = -R^f_{t,T}\mathrm{cov}_t\left[R_{t,T}, M_{t,T}\right] = R^f_{t,T}\beta_t\sigma^2_{W,t} \tag{11}$$

The variance of S&P 500 returns is $\sigma^2_{R,t} \equiv \beta^2_t\sigma^2_{w,t} + \sigma^2_{\varepsilon,t}$ and the Sharpe ratio is

$$\psi_t \equiv \frac{\mu_t - R^f_{t,T}}{\sigma_{R,t}} = \frac{R^f_{t,T}\beta_t\sigma^2_{W,t}}{\left(\beta^2_t\sigma^2_{W,t} + \sigma^2_{\varepsilon,t}\right)^{1/2}} \tag{12}$$

Using equation (5), we obtain $M^*_{t,T}$ by projecting the pricing kernel onto stock returns. We have

$$\begin{aligned} M^*_{t,T} &\equiv E_t\left[R^f_{t,T}M_{t,T}|R_{t,T}\right] \\ &= 1 - R^f_{t,T}E_t\left[W_T|R_{t,T}\right] + R^f_{t,T}\gamma_t E_t\left[\left(W^2_T - \sigma^2_{W,t}\right)|R_{t,T}\right] \\ &= 1 - \beta^M_t\left(R_{t,T} - \mu_t\right) + \gamma^M_t\left((R_{t,T} - \mu_t)^2 - \sigma^2_{R,t}\right) \end{aligned} \tag{13}$$

where $\beta^M_t = \psi_t\sigma^{-1}_{R,t}$ and $\gamma^M_t = \gamma_t\psi^2_t\left(R^f_{t,T}\sigma^2_{R,t}\right)^{-1}$. Online Appendix A.1 provides the full derivation (appendices are provided online as supporting information).

Changes in the two types of volatility, $\sigma_{W,t}$ and $\sigma_{\varepsilon,t}$, have different implications for $M^*_{t,T}$. An increase in idiosyncratic volatility ($\sigma_{\varepsilon,t}$) makes returns more noisy. Expected returns are unaffected by such an increase because systematic risk has not changed, but the additional noise in returns means that returns are less informative about the state variable. Thus an increase in $\sigma_{\varepsilon,t}$ shifts $M^*_{t,T}$ towards its unconditional mean of one. This result can be seen in equation (13) by noting that an increase in $\sigma_{\varepsilon,t}$ causes both β^M_t and γ^M_t to shrink towards zero—it causes the projected pricing kernel to become flatter.

In contrast, an increase in systematic volatility ($\sigma_{W,t}$) reduces the relative amount of noise in returns. Observed returns are now more informative about the state variable, and the projected pricing kernel steepens. In terms of equation (13), the derivatives of β^M_t and γ^M_t with respect to $\sigma_{W,t}$ are positive. Online Appendix A.2 provides additional details.

Using equation (13) along with the formula for the mean of a truncated normal distribution and some algebra provided in online Appendix A.1, we obtain

$$\mathrm{SPOCQ}_{t,T}(0,\theta) = \theta + \phi_\theta\psi_t\left(1 - q_\theta\frac{\gamma_t\psi_t}{R^f_{t,T}}\right) \tag{14}$$

where $q_\theta \equiv \Phi^{-1}(\theta)$ is the θ quantile, and $\phi_\theta \equiv \phi(q_\theta)$ is a constant equal to the standard normal density function evaluated at the θ quantile. Note that $\mathrm{SPOCQ}_{t,T}(0,1) = 1$.

In the absence of a risk premium, we have $\mathrm{SPOCQ}(0,\theta) = \theta$ for all θ. In the presence of a positive risk premium ($\beta_t > 0$), $\mathrm{SPOCQ}(0,\theta)$ exceeds θ for returns below the median. This result follows because $q_\theta \leq 0$ for any $\theta \leq 0.5$, and it arises in the model because the linear term in equation (13) implies that traders discount below-median returns more heavily than above-median returns.

If the stochastic discount factor has a quadratic component ($\gamma_t > 0$), then there exists a θ such that $\mathrm{SPOCQ}(0,\theta) < \theta$, i.e. one can always choose θ large enough such that $q_\theta\gamma_t\psi_t > R^f_{t,T}$. Thus, observing $\mathrm{SPOCQ}(\theta,1) > 1-\theta$ for some large θ implies that the pricing kernel is non-monotonic.

Increases in systematic volatility raise the left-tail SPOCQ, and increases in idiosyncratic volatility lower the left-tail SPOCQ. Specifically, in online Appendix A.2 we show that

$$\frac{\partial\mathrm{SPOCQ}_t(0,\theta)}{\partial\sigma_{W,t}} > 0 \quad \text{and} \quad \frac{\partial\mathrm{SPOCQ}_t(0,\theta)}{\partial\sigma_{\varepsilon,t}} < 0 \quad \text{for all} \quad \theta \leq 0.5 \tag{15}$$

DOI: 10.1002/jae

This result can also be seen by examining the Sharpe ratio (ψ_t). Systematic volatility increases the Sharpe ratio and therefore increases SPOCQ, whereas idiosyncratic volatility decreases the Sharpe ratio and decreases SPOCQ. If $\gamma_t = 0$, the sign of the derivatives is the same as in equation (15) for any θ; if $\gamma_t > 0$ the sign is ambiguous for $\theta > 0.5$.

In contrast to SPOCQ, evaluating the risk-neutral distribution at a fixed point R yields a more complicated expression that obfuscates changes in the pricing kernel:

$$F_t^*(R) = \Phi\left(\frac{R-\mu_t}{\sigma_{R,t}}\right) + \phi\left(\frac{R-\mu_t}{\sigma_{R,t}}\right)\psi_t\left(1 - \left(\frac{R-\mu_t}{\sigma_{R,t}}\right)\frac{\gamma_t\psi_t}{R_{t,T}^f}\right) \tag{16}$$

Overall, the relationship between $F_t^*(R)$ and the two types of volatility is ambiguous (see online Appendix A.2). However, for values of R in the left tail, an increase in volatility tends to increase $F_t^*(R)$, whether or not the additional volatility comes from an increase in systematic ($\sigma_{W,t}$) or idiosyncratic ($\sigma_{\varepsilon,t}$) volatility.

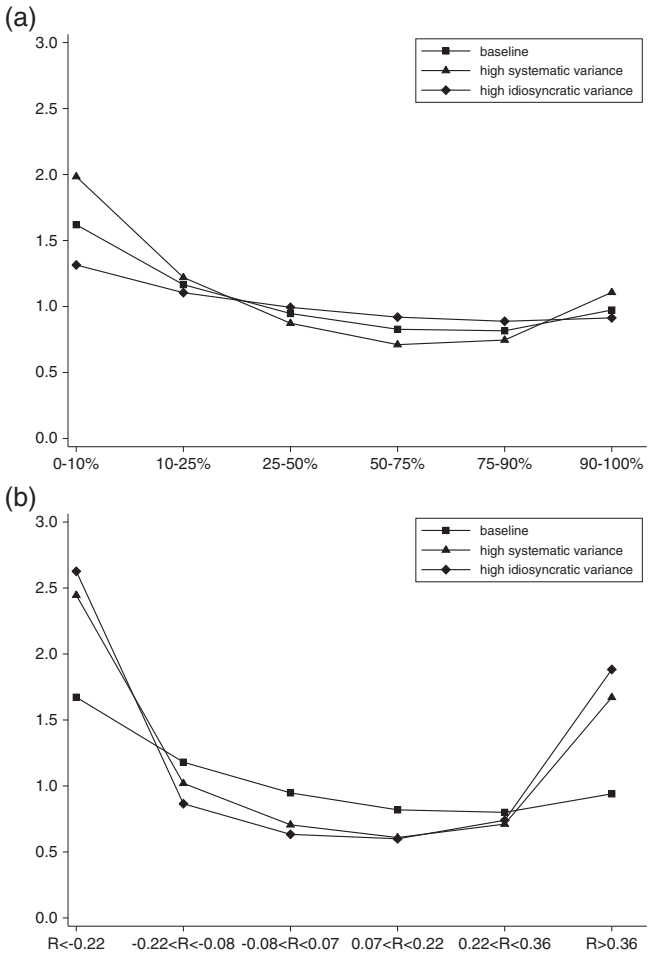

Figure 3. Projected pricing kernels from single-factor model: (a) quantile pricing kernel; (b) projected pricing kernel. We plot projected pricing kernels implied by the asset pricing model in Section 2.2. The baseline parameter settings are as follows: $R_{t,T}^f = 1.03$, $\sigma_{W,t} = 0.2$, $\sigma_{\varepsilon,t} = 0.1$, $\beta = 1$, and $\gamma = 4$. For high systematic variance, we set $\sigma_{W,t} = 0.26$. For high idiosyncratic variance, we set $\sigma_{\varepsilon,t} = 0.3$

DOI: 10.1002/jae

We illustrate these points in Figure 3, which plots the projected pricing kernel across several slices of the return distribution. Panel (a) shows the quantile pricing kernel (QPK):

$$\mathrm{QPK}_{t,T}(\theta_{j-1},\theta_j) \equiv \frac{\mathrm{SPOCQ}_{t,T}(\theta_{j-1},\theta_j)}{\theta_j - \theta_{j-1}} \tag{17}$$

for $\theta_j \in \{0, 0.1, 0.25, 0.5, 0.75, 0.9, 1\}$. When we increase the systematic volatility (σ_W) the left-tail QPK increases significantly. In contrast, the left-tail QPK decreases when we increase the idiosyncratic volatility (σ_ε). The change in the right-tail QPK is ambiguous in general, and is approximately zero for the parameter values in Figure 3.

Panel (b) shows the projected pricing kernel across fixed return ranges given by

$$\mathrm{PPK}_{t,T}(R_{j-1},R_j) = \frac{F_t^*\left(R_j\right) - F_t^*\left(R_{j-1}\right)}{a_j - a_{j-1}} \tag{18}$$

for $R_j \in \{-\infty, 0.78, 0.92, 1.07, 1.22, 1.36, \infty\}$ and $a_j \in \{0, 0.1, 0.25, 0.5, 0.75, 0.9, 1\}$. We divide by $a_j - a_{j-1}$ to put both panels in the same scale. Panel (b) is calibrated so that the baseline is the same as in panel (a), i.e. we chose the values R_j such that $\Pr(R_{t,T} \leq R_j) = a_j$ with the baseline parameters. Increases in volatility of either source raise the PPK in both tails because an increase in volatility increases the probability that a return exceeds the fixed thresholds. In contrast, SPOCQ and the QPK reveal changes in the projected pricing kernel unconfounded by changes in the physical return distribution because they hold constant the probability of each event.

3. ESTIMATING THE COMPONENTS OF SPOCQ

To obtain SPOCQ, we need estimates of the risk-neutral distribution function $F_t^*(\cdot)$ and the conditional quantile $q_t(\theta)$ at which to evaluate this function. In Section 3.1 we describe how we estimate $F_t^*(\cdot)$ from the cross-section of options prices. We present our approach for estimating $q_t(\theta)$ in Section 3.2.

3.1. Risk-Neutral Distribution Implied by Options Prices

A large literature has developed methods for estimating the risk-neutral distribution directly from options prices (e.g. Breeden and Litzenberger, 1978). Our approach relies heavily on this literature. Using X and S_T to denote the strike and underlying price at the expiration date T, the price of a European put option may be written as

$$P_t(X,T) = E_t\left[M_{t,T} \times \max(X - S_T, 0)\right] = E_t\left[M_{t,T}\right]\int_{-\infty}^{X}(X - S_T)\mathrm{d}F_t^*(S_T) \tag{19}$$

We define the adjusted put option price, also known as the forward option price, as

$$\tilde{P}_t(X,T) \equiv \frac{P_t(X,T)}{E_t\left[M_{t,T}\right]} \tag{20}$$

and take the derivative with respect to the strike price to get

$$\frac{\partial \tilde{P}_t(X,T)}{\partial X} = \int_{-\infty}^{X}\mathrm{d}F_t^*(S_T) = F_t^*(X) \tag{21}$$

DOI: 10.1002/jae

Put-call parity produces a parallel expression for the adjusted call price as

$$\tilde{C}_t(X,T) \equiv \frac{C_t(X,T) - S_t + D_{t,T}}{E_t\left[M_{t,T}\right]} + X \tag{22}$$

where $D_{t,T}$ denotes the present value of dividends to be paid before T. Differentiating yields

$$\frac{\partial \tilde{C}_t(X,T)}{\partial X} = F_t^*(X) \tag{23}$$

We obtain the risk-neutral distribution $F_t^*(S_T)$ by estimating the first derivative of the adjusted call and put option price curves, $\tilde{C}_t(X,T)$ and $\tilde{P}_t(X,T)$, with respect to X.

We use a mixture of logistic distributions to approximate the risk-neutral distribution of the adjusted option prices in equations (20) and (22). A mixture of distributions offers a flexible approximation to any distribution (Marron and Wand, 1992). By using a parametric distribution, we avoid the problems that other curve-fitting methods (e.g. splines) have in estimating the tails, and getting the estimated distribution to integrate to one (see Figlewski, 2010). The logistic distribution is appealing because its integral exists in closed form, which enables us to work with the observed options pricing curve directly.

Instead of fitting a distribution to the derivatives of the adjusted option prices, we fit the integral of a distribution to the adjusted options prices themselves. Fitting the curve before differentiating the option pricing curve avoids arbitrary assignment of the point at which the derivative applies. Thus, by using a mixture of logistic distributions, we are able to fit a flexible function to the adjusted option prices, and simultaneously impose the restriction that the derivative is a distribution.

Using X_T to indicate the exercise price of an option that expires at T, we fit the function

$$F_t^*(X_T) = \sum_{j=1}^{J} \omega_{jt} \Lambda_{jt}\left(X_T; \mu_{jt}, \sigma_{jt}\right) \tag{24}$$

where $\Lambda_{jt}(\cdot)$ is the logistic distribution. We do so by fitting the options pricing curve to the integral of $F_t^*(X_T)$ and specifying the model

$$o_{it} = \sum_{j=1}^{J} \omega_{jt} \widetilde{\Lambda}_{jt}\left(X_{iT}; \mu_{jt}, \sigma_{jt}\right) + \varepsilon_{it} \tag{25}$$

where o_{it} denotes the ith adjusted options price that we observe, X_{iT} is the strike price of that option, and

$$\widetilde{\Lambda}_{jt}\left(X_T; \mu_{jt}, \sigma_{jt}\right) = \int \Lambda_{jt} \mathrm{d}X_T = \sigma_{jt} \ln\left(\exp\left(\frac{X_T}{\sigma_{jt}}\right) + \exp\left(\frac{\mu_{jt}}{\sigma_{jt}}\right)\right) \tag{26}$$

We impose $\sum_{j=1}^{J} \omega_{jt} = 1$, with $\omega_{jt} \geq 0$, $j = 1, \ldots, J$.[1] We fit the model in equation (25) separately for each trading date t. In our application we set $J = 2$, so we fit the distribution using five parameters. We repeated our analysis with $J = 3$ and the results are highly similar.

This number of parameters provides considerable flexibility when compared to parametric distributions, such as the normal (two parameters), skew-normal and Student's t (three parameters), as well as

[1] When computing the adjusted call price, we set $D_{t,T} = 0$. We then include a dummy variable in equation (25) that equals one for call options and zero for puts. This dummy adjusts for the omitted dividends, which are constant across strike prices. See online Appendix A.3 for details.

DOI: 10.1002/jae

skew-t (four parameters; see Azzalini and Capitanio, 2003). Online Appendix A.3 provides details on how we estimate the parameters in equation (26) using constrained nonlinear least squares to produce the estimated risk-neutral distribution:

$$\widehat{F}_t^*(X_T) = \sum_{j=1}^{J} \widehat{\omega}_{jt} \Lambda_{jt}\left(X_T; \widehat{\mu}_{jt}, \widehat{\sigma}_{jt}\right) \tag{27}$$

3.2. Conditional Quantiles

For each trading date t, we require quantiles of the return a trader could earn by holding the asset until T. Existing models for estimating conditional quantiles include quantile regression (Koenker and Bassett, 1978), univariate CAViaR models (Engle and Manganelli, 2004), multi-quantile CAViaR models (White *et al.*, 2010) and fully parametric generalized autoregressive conditional heteroskedasticity (GARCH) models (e.g. Christoffersen *et al.*, 2013). Among these candidate models, the best may depend on the details of a particular application, especially on the sequence of dates t and T.

The S&P 500 options have one expiration date per month; hence we use a horizon of 28 calendar days in our empirical analysis ($T - t = 28$). We therefore have a sequence of monthly observations on a 28-day return, which implies a time series of non-overlapping returns. S&P 500 option payouts are determined by the value of the index at the market open on the third Friday of the month. We are interested specifically in returns from the close of business 4 weeks prior to expiration (a Friday) until the open 4 weeks ahead. Thus we model the conditional quantiles of the log return series for time $t = 1, \ldots, T$ defined by

$$r_{t,T} \equiv 100 \times \ln\left(\mathrm{SP}_T^{\mathrm{open}} / \mathrm{SP}_t^{\mathrm{close}}\right) \tag{28}$$

where SP_t denotes the S&P 500 index value at t, and T occurs 28 calendar days after t.

For our baseline analysis, we use a quantile-regression model to estimate conditional quantiles. We tried several CAViaR models, but found that they did not pass diagnostic tests so we discarded them.[2] We also estimated the discrete analog of the Heston and Nandi (2000) GARCH in Christoffersen *et al.* (2013) and used it to project the distribution of the return over the ensuing 28 days at each date t. However, we found that the quantiles generated by this method were too smooth, failing both conditional and unconditional diagnostic tests. Specifically, this GARCH model severely underestimated the probability of tail events, especially when volatility was high.

We use volatility as the main explanatory variable in our quantile regressions. Specifically, we use the square root of *realized* continuous variation of daily returns for the S&P index over the last 20 trading days (CV) and the Chicago Board Options Exchange (CBOE) Volatility Index (VIX).[3] The VIX includes a forward-looking component and a volatility-risk premium component, neither of which are in realized volatility. The first component is relevant for forecasting quantiles, but the second is not. We allow the data to determine whether adding VIX improves the quantile predictions, and we parametrize the model so that the coefficients directly reveal the incremental effect of adding VIX. The quantile regression model is

[2] Our biggest challenge with the CAViaR models was modeling the left tail of the distribution and, in particular, the 25% quantile. The equation for the 25% quantile gave us an autoregressive parameter very close to one, which generated a very smooth quantile series and hit rates that were well below the nominal level of 25%. We also experienced issues with the 50% quantile, such as low hit rates.

[3] Our measure of continuous variation follows Bollerslev and Todorov (2011) and is described in detail in online Appendix A.6. We use realized continuous variation rather than realized total variation because we found that removing jumps significantly improved the fit of the model.

DOI: 10.1002/jae

$$r_{t,T} = \mathbf{x_t}'\beta_\theta + \varepsilon^\theta_{t,T} \tag{29}$$

and the conditional quantiles are $q_t(\theta) = \mathbf{x_t}'\beta_\theta$, where $\mathbf{x_t}' = (1, \mathrm{CV}_t, \mathrm{VIX}_t - \mathrm{CV}_t)$. We also estimated quantile regression models that included other explanatory variables computed over the past 20 trading days, such as empirical quantiles, estimated jump tails as in Bollerslev and Todorov (2011) and lagged index returns. These variables often failed to be statistically significant and including them did not change our results.

We assess the fit of our conditional-quantile models using a series of backtests that have been employed to evaluate value-at-risk (VaR) model performance. If a conditional-quantile model is correctly specified, then hits[4] $h_t(\theta)$ occur with conditional probability θ, and thus no variable should be able to predict hits. Thus the backtests are designed to test the null hypothesis that hits are not predictable. We provide details in online Appendix A.4.

The SPOCQ itself also provides a way to evaluate our quantile estimates. If the conditional-quantile model is correctly specified, then SPOCQ represents the value of a dollar in the event that $r_{t,T}$ lies below the $q_t(\theta)$ quantile of the distribution—an event with probability θ. If the quantile model is misspecified, then the estimated SPOCQ represents an event with a different probability than θ. For example, if the probability that $r_{t,T} \leq q_t(\theta)$ exceeds θ, then our estimated SPOCQ will be too large. This large SPOCQ value reflects a high *probability* rather than a high *value* for the event under consideration. Therefore, we test the following moment condition implied by the lack of hit predictability:

$$E\left[\widehat{F}^*_t\left(\mathbf{x_t}'\beta_\theta\right)\left(h_t(\theta) - \theta\right)\right] = 0 \tag{30}$$

If the quantile regression model is correctly specified, then we should not be able to reject the null hypothesis that equation (30) is a valid over-identifying moment condition using a J-statistic.

3.3. Inference

The following decomposition reveals two sources of estimation error in the estimated SPOCQ:

$$\widehat{\mathrm{SPOCQ}}_{t,T}(0,\theta_j) = \widehat{F}^*_t\left(\mathbf{x_t}'\widehat{\beta}_{\theta_j}\right) = \mathrm{SPOCQ}_{t,T}(0,\theta_j) + u_t + v_t \tag{31}$$

where

$$\begin{aligned} u_t &= \widehat{F}^*_t\left(\mathbf{x_t}'\beta_{\theta_j}\right) - F^*_t\left(\mathbf{x_t}'\beta_{\theta_j}\right) \\ v_t &= \widehat{F}^*_t\left(\mathbf{x_t}'\widehat{\beta}_{\theta_j}\right) - \widehat{F}^*_t\left(\mathbf{x_t}'\beta_{\theta_j}\right) = \widehat{f}^*_t\left(\mathbf{x_t}'\beta_{\theta_j}\right)\mathbf{x_t}'\left(\widehat{\beta}_{\theta_j} - \beta_{\theta_j}\right) + o_p(1) \end{aligned}$$

and f^* denotes the risk-neutral PDF. The last line comes from a first-order expansion; the approximation error is $o_p(1)$ because $\widehat{\beta}_{\theta_j}$ is consistent for β_{θ_j}.[5]

We use u_t to denote the error from using the logistic-mixture model to approximate the risk-neutral distribution, which is indicated by the hat over F in equation (31). We use v_t to denote the estimation error from the quantile regression, which is indicated by the hat over β_θ in equation (31). We fit the logistic mixture separately for each time period, so the first type of error is particular to each time

[4] The hits variable equals one if the return is less than the conditional quantile and zero otherwise.

[5] The analysis in this section also applies to $\widehat{\mathrm{SPOCQ}}_{t,T}(\theta_{j-1},\theta_j)$ and to the quantile pricing kernel $\widehat{\mathrm{QPK}}_{t,T}(\theta_{j-1},\theta_j) \equiv \widehat{\mathrm{SPOCQ}}_{t,T}(\theta_{j-1},\theta_j)/(\theta_j - \theta_{j-1})$. For clarity and brevity, we discuss only the $\widehat{\mathrm{SPOCQ}}_{t,T}(\mathbf{0},\theta_j)$ case here.

DOI: 10.1002/jae

period. In contrast, the second type of estimation error arises in repeated samples; an error in $\widehat{\beta}_\theta$ affects the estimated SPOCQ for all t in a given sample.

Later, we compute averages of SPOCQ over time and we regress our estimated SPOCQ on a set of covariates. Including the two sources of estimation error, the regression equation is

$$\widehat{\text{SPOCQ}}(0,\theta_j) = Z\gamma + \eta \tag{32}$$

where $\widehat{\text{SPOCQ}}(0,\theta_j)$ denotes an $n\times 1$ vector of SPOCQs, Z is a matrix of covariates and $\eta = \varepsilon + u + v$ is an $n \times 1$ error vector. The error vector is non-spherical due to possible autocorrelation in ε and u and to the quantile regression error v. We assume the logistic-mixture estimation error averages zero and is uncorrelated with Z. The quantile regression error is asymptotically uncorrelated with Z because of the moment condition underlying $\widehat{\beta}_\theta$.[6] It follows that the OLS estimator is consistent for γ.

To account for the non-spherical errors—in the spirit of Chamberlain (1982) and Nevo (2001)—we estimate γ by generalized least squares (GLS). The error variance is

$$\text{var}[\eta|Z,X] \equiv \Sigma = \text{var}[\varepsilon + u|Z,X] + f^* X V_{\widehat{\beta}_\theta} X' f^* + o_p(1) \tag{33}$$

where $V_{\widehat{\beta}_\theta} \equiv \text{var}[\widehat{\beta}_\theta]$ and f^* denotes a diagonal matrix with each element equal to the risk-neutral PDF for the corresponding observation. We estimate the first term in equation (33) by specifying an AR(1) process for η_t. For the second term, we use a robust estimate of $V_{\widehat{\beta}_\theta}$ from the quantile regression.

3.4. Data

The data for the S&P 500 are from the Commodity Research Bureau and span January 1990 to April 2012. The mean monthly return during this period was 0.6%, and the standard deviation was 4.6%. The data for the S&P 500 options are from the CBOE. The market for S&P 500 options operates between 8:30 am and 3:15 pm central time. At any point in time, three near-term expiration months are trading along with three additional months from the March quarterly cycle (March, June, September and December). Currently, the strike price intervals are set at 5 points (25-point intervals for distant expiration months).

S&P 500 options expire on the Saturday following the third Friday of the month. The exercise-settlement value equals the opening value of the index on the last business day (usually a Friday) before the expiration date, and trading typically ceases one day earlier (usually a Thursday). The options may be exercised only on the last business day before expiration (European style), with exercise resulting in delivery of cash on the business day following expiration. We also observe daily trading volume, open interest, as well as closing bid and ask quotes, for calls and puts between January 1990 and April 2012.

We measure the option price using the closing mid quote, defined as the average of the bid and ask quotes, on the Friday 4 weeks prior to expiration. Following Figlewski (2010), we drop observations with mid quotes below 0.5 and those with no trading volume on day t.[7] We calculate the adjusted options prices in equations (20) and (22), using $E_t\left[M_{t,T}\right]^{-1} = (1 + b_{t,T})$, where $b_{t,T}$ denotes the LIBOR rate that covers the interval from t to T.[8]

[6] The asymptotic first-order condition for quantile regression is $\sum_{t=1}^n x_t(\theta - h_t(\theta)) = 0 + o_p(1)$, so a first-order expansion implies $\widehat{\beta}_\theta - \beta_\theta = -\left(\sum_{t=1}^n x_t x_t' f_t\right)^{-1}\left(\sum_{t=1}^n x_t(\theta - h_t(\theta))\right) + o_p(1)$; see, for example, Buchinsky (1998). Thus $E[z_t v_t] = -E\left[z_t \widehat{f}_t^*\left(\mathbf{x_t}'\beta_{\theta_j}\right)\mathbf{x_t}'\left(\sum_{t=1}^n x_t x_t' f_t\right)^{-1}\sum_{t=1}^n x_t(\theta - h_t(\theta))\right] + o_p(1) = 0 + o_p(1)$.

[7] We exclude 'mini' contracts identified by the following codes 'SXZ', 'SPB', 'LSW', 'LSX', 'LSY', 'LSZ', 'XSC', 'XSE', 'XSK', 'XSL', 'XSO','XSP'.

[8] We use the 1-month London Interbank Offered Rate (LIBOR) series from FRED.

DOI: 10.1002/jae

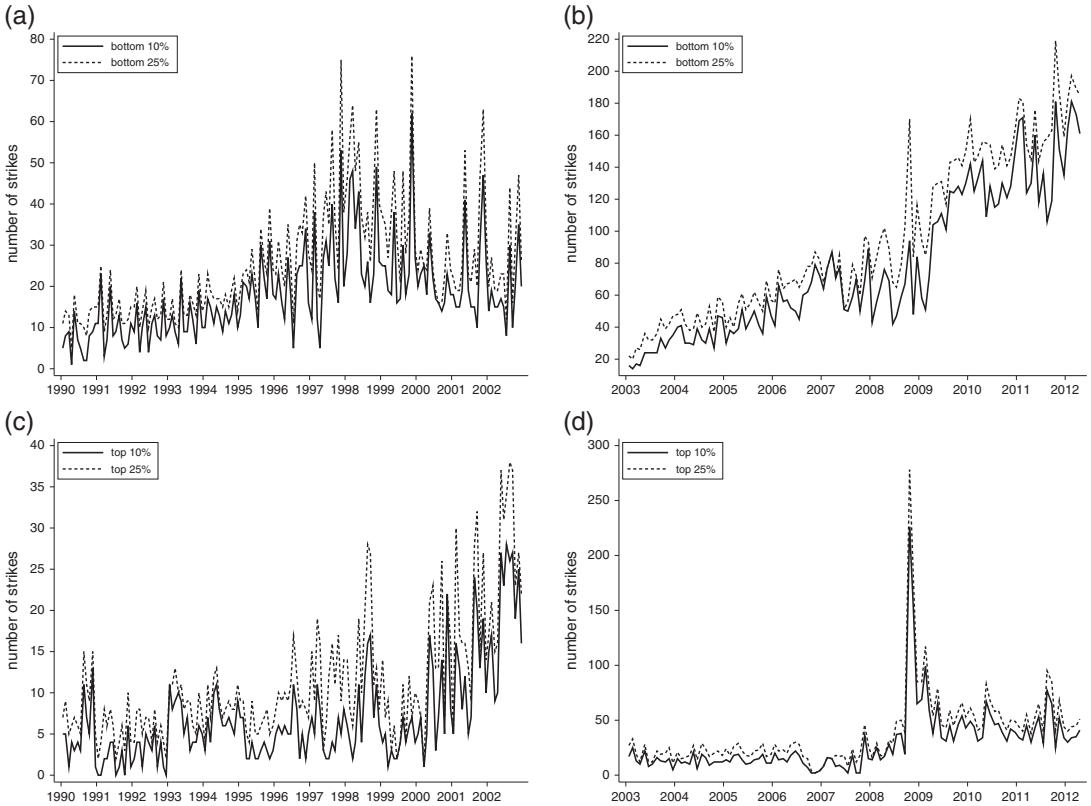

Figure 4. Number of strikes in various parts of the risk-neutral distribution: (a) left tail, 1990–2002; (b) left tail, 2003–2012; (c) right tail, 1990–2002; (d) right tail, 2003–2012. We plot the total number of strikes that falls within various parts of the risk-neutral distribution using its 10% 25%, 75% and 90% quantiles. Panels (a) and (c) have different scales on the vertical axis compared to panels (b) and (d) because of the large increase in the number of strikes in the early 2000s.

4. ESTIMATION RESULTS

We fit the model in equation (25) for each of the 268 trading dates in our full sample, calculating the adjusted options prices in equations (21) and (22) using the LIBOR rate that covers the interval from t to T as our measure of $R^f_{t,T}$.[9] Figure 4 contains the number of strikes for both call and put options for various slices of the risk-neutral distribution of the S&P 500 on each trading date. The number of strike prices increases throughout the sample, which is consistent with the increasing liquidity of the S&P 500 derivatives market over time. Although the number of strikes is low in some months early in the sample, we always have at least a few in each of the top and bottom deciles.[10]

The most likely component in the mixture has an average weight of $\omega_t = 0.8$ across the 268 dates. The Bayesian information criterion (BIC) favors the single-component model over the mixture model for all trading dates except a few in 1993–1994. Because estimates of a mixing model remain consistent if the specified number of mixture components exceeds the true number of components (Cho

[9] To check the robustness of this measure, we repeated our analysis treating $\boldsymbol{R}^f_{t,T}$ as a parameter to be estimated when we fit the risk-neutral distribution to option prices. This alternative approach does not change our results, so we proceed with LIBOR. Online Appendix Figure A.1 provides a plot of the weight, scale and location parameters associated with the most likely component of the mixture for each of the 268 trading dates in our sample.

[10] The prominent spike in panel (d) of the figure corresponds to October 2008—the peak of the most recent financial crisis.

DOI: 10.1002/jae

and White, 2007), we proceed with the two-component model for all dates. Our results are unchanged if we instead use a three-component model or if we choose each period the model selected by BIC.

Table I presents quantile regression estimates along with backtests for each of the five conditional-quantile models we considered. We estimate each of the five quantiles using 268 monthly observations between January 1990 and April 2012. Each of the dates in these monthly series correspond to the trading dates of our options data.

Continuous volatility (CV) enters with a negative sign in the equations for the first two quantiles and with a positive sign for the remaining quantiles. It is significant at 1% for the 90% quantile. It is also significant at 5% for the 10% and 75% quantiles, and at 10% for the 25% quantile. The coefficient on VIX-CV is significant and large only for the two right-tail quantiles. Hansen's J-statistic calculated using equation (30) rejects the null of correct over-identifying restrictions for the 90% quantile, suggesting that this quantile may not be well specified. The DQ and VQR tests reinforce this finding, but most of the other backtests do not reject the model—see online Appendix A.4 for details on the tests. An exception is the VQR test, which rejects the model for the 75% quantile.

If our quantile model is well specified, then we should not be able to predict whether the return is likely to exceed the quantile. We find evidence of predictability only for the 90% quantile. To the extent that our conditional quantiles are incorrect, the resulting SPOCQ series cannot be interpreted

Table I. Quantile regression: estimates and model diagnostics

	Conditional quantile				
Variable	10%	25%	50%	75%	90%
Panel A: Estimates					
CV	−5.3548**	−2.1562*	0.5649	1.7149**	4.3788***
	(2.6450)	(1.1951)	(0.9053)	(0.8419)	(0.9154)
VIX-CV	−0.2768	1.0831	1.0581	3.2634**	3.7462***
	(3.7861)	(1.5601)	(1.3313)	(1.4250)	(0.7009)
Constant	0.0023	0.0008	0.0052	0.0081	0.0013
	(0.0221)	(0.0101)	(0.0081)	(0.0078)	(0.0078)
Panel B: Measures of fit					
J-statistic	0.330	0.390	0.408	0.307	4.775**
	[0.566]	[0.532]	[0.523]	[0.579]	[0.029]
LM test	8.751**	3.606	1.759	8.041**	11.646***
	[0.013]	[0.165]	[0.415]	[0.018]	[0.003]
Hits	10%	25%	50%	74%	91%
R^2	0.069	0.017	0.007	0.055	0.141
Panel C: Diagnostic tests for model fit					
DQ test	0.423	0.618	0.336	0.539	5.135**
	[0.515]	[0.432]	[0.562]	[0.463]	[0.024]
VQR test	0.524	0.424	1.070	16.060***	10.460***
	[0.769]	[0.809]	[0.586]	[0.000]	[0.005]
LR IND test	4.686	2.834	2.009	0.668	0.407
	[0.970]	[0.908]	[0.844]	[0.586]	[0.477]
LR CC test	4.688	2.834	2.024	0.747	0.544
	[0.904]	[0.758]	[0.636]	[0.312]	[0.238]
Bartlett's test	0.666	0.685	0.938	0.703	0.663
	[0.767]	[0.737]	[0.342]	[0.707]	[0.771]
Portmanteau test	11.120	14.760	14.400	15.110	9.052
	[0.519]	[0.255]	[0.276]	[0.235]	[0.698]
Obs.	268	268	268	268	268

Note: We report estimates and model diagnostics for the quantile-regression model in (29). Standard errors are in parentheses and are robust to heteroskedasticity. We report p-values in squared brackets. The asterisks indicate statistical significance as follows: 1% (***), 5% (**), 10% (*). The details of the tests in Panel C are available in Appendix A.4.

DOI: 10.1002/jae

as the state price of a set of events with known and constant probability. In online Appendix A.5, we show robustness to alternative quantile specifications.

5. SPOCQ FOR S&P 500 RETURNS

In this section, we discuss the most salient SPOCQ features during the period 1990–2012. We estimate SPOCQ each month by evaluating the risk-neutral distribution function computed as in Section 3.1 at the estimated conditional quantiles computed as in Section 3.2. Panel (a) of Figure 5 shows the five conditional quantiles we obtained using the estimates in Table I, and panel (b) shows the implied SPOCQ series.

5.1. Risk Aversion

Under risk neutrality, the five SPOCQ series would coincide with the straight lines at 0.10, 0.25, 0.50, 0.75 and 0.90, respectively. As Figure 5(b) shows, the SPOCQ exceeds its risk-neutral counterpart in almost all months for quantiles at the median and below. Thus investors are willing to pay a premium for a dollar in the event that the S&P 500 produces a negative return. This premium averages 6 cents for the bottom decile, 7 cents for the bottom quartile and 9 cents for below-median returns. In terms of the model in Section 2.2, this finding is consistent with a positive equity risk premium on average ($\beta_t > 0$).

5.2. Noise

The SPOCQ series in Figure 5 exhibit substantial month-to-month variation. The series for the four lower quantiles each have standard deviations between 0.04 and 0.06 and exhibit month-to-month changes that frequently exceed 0.05. Conditional on the estimated risk-neutral distribution, there is a one-to-one mapping between SPOCQ and the conditional quantiles. Therefore, *less* variation in the estimated quantile series implies *more* variation in SPOCQ series and vice versa.

If the true conditional quantiles contain a component that is not predictable based on lagged observables, our estimated quantile series would be too smooth and the implied SPOCQ series would be too noisy. In such a case, SPOCQ would have power to predict quantile hits. The J-statistic in Table I shows that SPOCQ does not significantly predict quantile hits for the four lower quantiles. It shows significant predictive power only for the 90% quantile; for all other quantiles the p-value exceeds 0.5.

To further investigate whether the noise in SPOCQ reflects quantile estimation error, we extract the noise component and check whether it can predict returns. If the noise predicts returns with a negative coefficient, then it suggests that a high SPOCQ in a given month indicates that a low return is likely rather than a high discount rate for low returns.[11] We fit ARMA(1,1) models to each of the five SPOCQ series, and we regress $r_{t,T}$ on the ARMA(1,1) residuals.[12] In univariate regressions, the coefficient on the noise has inconsistent signs across regressions and is statistically insignificant in all cases with t-statistics less than one for all except the 90% quantile. In a multivariate regression including the residuals from each quantile, the F-statistic for joint significance has a p-value of 0.77 and none of the five SPOCQ noise series has a significant t-statistic. We repeated this exercise with quantile hits as the left-hand-side variable and obtained the same results—weak evidence of predictability for the 90% quantile and none for the other quantiles.

[11] Suppose the quantile is wrong, so that a high SPOCQ indicates a high event probability rather than a high state price. If this is true, then SPOCQ should predict returns. The coefficient should be negative because a high state probability means a high chance of a return less than q.

[12] See Table A.2 in the online Appendix.

DOI: 10.1002/jae

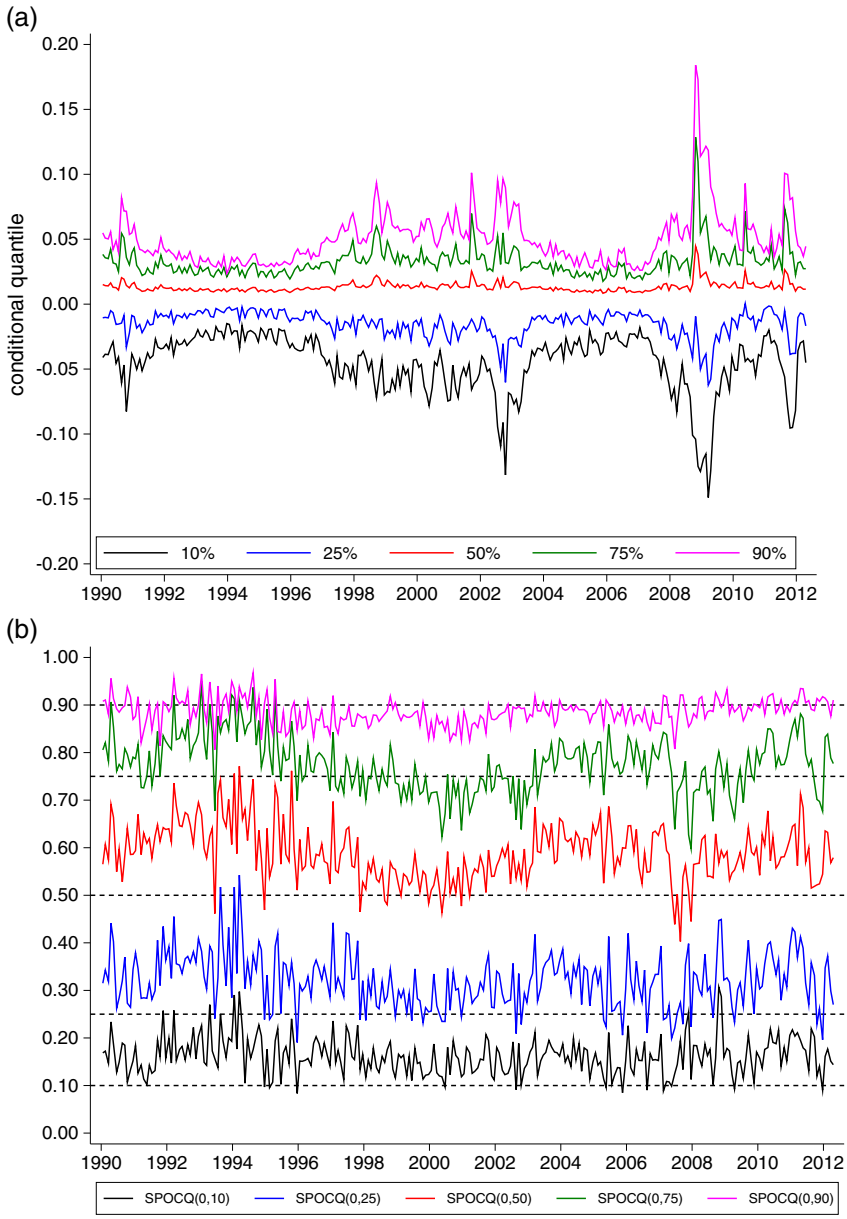

Figure 5. Conditional quantiles and SPOCQ: (a) quantiles; (b) SPOCQ. Panel (a) shows the monthly time series of the conditional quantiles. Panel (b) shows the monthly time series of the implied SPOCQ. Each monthly observation corresponds to a trading date in our options data. For brevity, we write quantiles as integers. This figure is available in color online at wileyonlinelibrary.com/journal/jae

Estimation error in the risk-neutral distribution would also generate estimation error in SPOCQ. Such error could arise from insufficient strike prices being available in a particular interquantile range. Figure 4 shows that this could be the case for some months early in the sample. Estimation error could also arise from the fact that we use mid quotes to proxy for options prices. If the true options pricing curve is closer to the ask in some months and closer to the bid in other months, then using mid quotes could induce noise in SPOCQ. To address this possibility, we re-estimated our SPOCQ series using

DOI: 10.1002/jae

only bid quotes and using only ask quotes. Compared to those from mid quotes, the resulting SPOCQ series were very similar for the bids and slightly more volatile for the ask. Estimation error could also arise if the logistic mixture is not flexible enough to fit the risk-neutral distribution. However, our results do not change when we increase the number of terms in the mixture to 3.

Overall, the substantial high-frequency variation in the projected pricing kernel is not attributable to estimation error in the quantiles or the risk-neutral distribution, especially after the first few years of the sample.

5.3. Financial Crises

World financial markets experienced a sequence of financial crises between 2007 and the end of our sample. The US credit crunch began in August 2007 and intensified in September 2008 after the Lehman Brothers bankruptcy. Sovereign debt crises then hit several European countries in 2010 and 2011. The USA and many other countries experienced deep recessions during this period, which suggests a substantial increase in systematic risk.

Figure 6 plots the bottom-decile and the bottom-quartile SPOCQ from 2006 through the end of the sample. We standardize by the probability that returns fall in the relevant interquantile range so that each element in the figure is measured on the same scale; that is, we use the quantile pricing kernel (QPK) defined in equation (17). We plot quarterly averages to smooth out the noise. Both series reach their peak in the fourth quarter of 2008, which is the period immediately after the Lehman Brothers bankruptcy. Recalling from Section 2.2 that an increase in systematic risk increases left-tail SPOCQ, this finding is consistent with the fourth quarter of 2008 being a period of very high systematic risk. However, both series dropped back below sample-average levels in the first quarter of 2009. The S&P 500 bottomed out during the first quarter of 2009, and the figure indicates relatively low systematic risk throughout that calendar year.

The three next largest values for the bottom decile QPK occurred in 2011:Q1, 2010:Q2 and 2007:Q4. The first two of these dates also coincide with peaks in the bottom-quartile QPK. The high values of left-tail QPK in 2010–11 suggest that the market viewed the European sovereign debt crises as significant systematic events. The peak in 2007:Q4 coincides with the beginning of the US credit

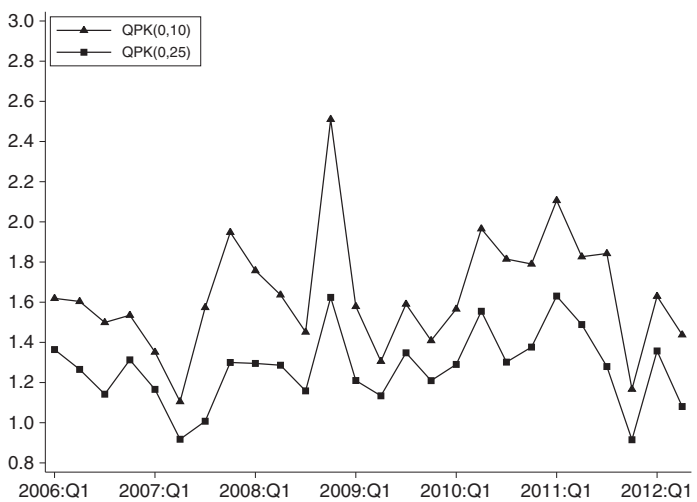

Figure 6. Left-tail quantile pricing kernel 2006:Q1–2012:Q1. We plot quarterly averages of the bottom-decile (QPK(0,10)) and the bottom quartile (QPK(0,25)) quantile pricing kernels using the expression in equation (17). For brevity, we write quantiles as integers

DOI: 10.1002/jae

crunch and the NBER-dated business cycle peak. Overall, the behavior of the left-tail QPK in 2006–12 shows that the major financial events in that period coincided with a high estimated price of risk.

The left-tail SPOCQ values during the financial crisis are similar to the values observed in the mid 1990s (see Figure 5). Based on the model in Section 2.2, this result implies little idiosyncratic (i.e. equity-specific) volatility in the mid 1990s and relatively more idiosyncratic volatility during the financial crisis. Macroeconomic uncertainty was higher in 2008 than in the mid 1990s, but equity-specific volatility may also have been higher. Thus S&P 500 returns would have been relatively less informative about the state of the macroeconomy in 2008 than in the mid 1990s when equity markets were calm.

5.4. U-Shaped Pricing Kernel

The SPOCQ series for above-median returns are generally *inconsistent* with standard notions of risk aversion. They show evidence of an upward-sloping part in the pricing kernel consistent with the pricing kernel puzzle. Given that SPOCQ(0,θ)$> \theta$ for the left-tail series, a finding of SPOCQ(0,θ)$< \theta$ in the right tail is sufficient for the pricing kernel to be upward sloping. The SPOCQ(0,90) series is less than 0.9 for about 70% of the sample, especially during the highly volatile period between 1996 and 2004. The SPOCQ(0,75) series exhibits more variation than SPOCQ(0,90), and it falls below 0.75 for a third of the sample.

To highlight changes in pricing kernel over time, Figure 7 shows the average QPK over several subsamples. Under risk neutrality, the QPK would equal one for all quantiles; values greater than one indicate a larger than average discount rate for that quantile range. With the exception of 2008–12, each curve in Figure 7 is U-shaped, which illustrates the pricing kernel puzzle. In all subperiods, QPK values in the left tail of the distribution exceed those in the right tail, so large negative returns are discounted more than large positive returns, which is consistent with standard theory.

Panel (f) shows the QPK for the months with low and high volatility.[13] The pricing kernel puzzle was most prominent in the years that exhibited the highest volatility in the sample. The average top-quartile QPK over the whole sample was 0.90, and the average for the third quartile was 0.74. Figure 1 also illustrates the positive correlation between the right-tail QPK and volatility. The derivations in Section 2.2 show that we cannot deduce from right-tail QPK behavior whether volatility changes are systematic or idiosyncratic. We investigate potential explanations for the pricing kernel puzzle in Section 6.3.

Our U-shaped pricing kernel finding is consistent with Jackwerth (2000) and, more recently, with Bakshi *et al.* (2010). Song and Xiu (2016) project the pricing kernel onto both returns and volatility, and find that the pricing kernel puzzle dissipates. They show that upward-sloping portion of their pricing kernel is imprecisely estimated and therefore could be an artifact of statistical error. An interesting topic for further research is whether a two-dimensional projection would also attenuate the pricing kernel puzzle in our setting.

5.5. Volatility: Systematic or Idiosyncratic?

Across our sample period, the bottom-decile QPK exhibits a correlation of -0.40 with volatility. Figure 1 shows a similar negative correlation for the bottom-quartile QPK. Based on our arguments in Section 2.2, this finding implies that most variation in volatility comes from changes in idiosyncratic volatility—risk for which investors do not require compensation.[14]

[13] Months during which volatility is below (above) its 1st (3rd) quartile are treated as low- (high-)volatility months. Recall that there are 268 months in our sample. Hence there are 67 low and 67 high-volatility months.

[14] We also computed state prices at unconditional quantiles as in equation (16), i.e. we evaluated the risk-neutral distribution at the same point each month. Consistent with the predictions of Section 2.2, the resulting state prices were positively correlated with volatility in both tails.

DOI: 10.1002/jae

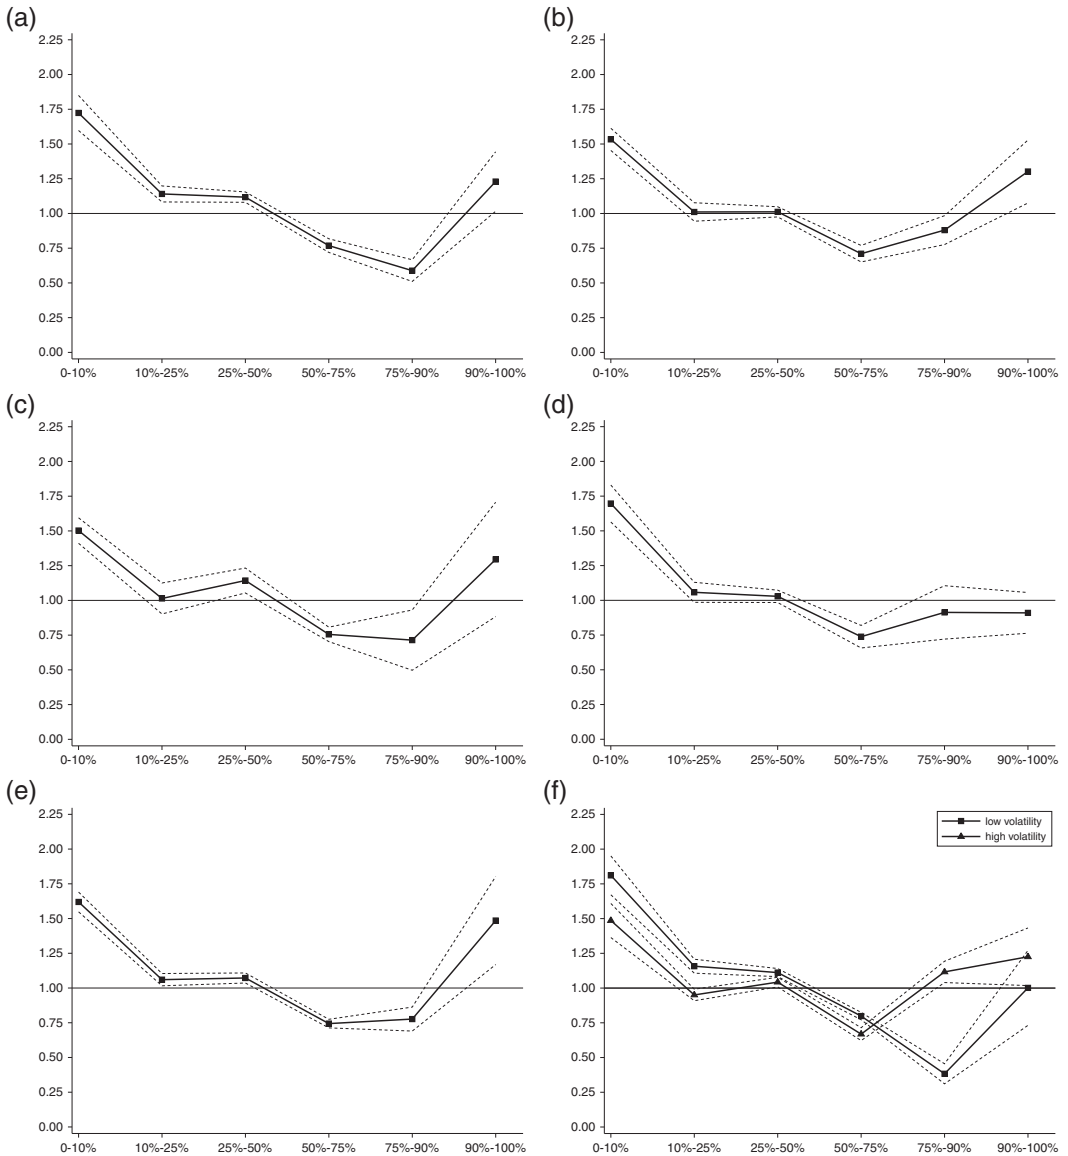

Figure 7. Averaging the quantile pricing kernel: (a) 1990–1996; (b) 1997–2003; (c) 2004–2007; (d) 2008–2012; (e) 1990–2012; (f) by volatility period. Each plot in the figure shows an average quantile pricing kernel estimated by GLS as in Section 3.3. For brevity, we write quantiles as integers. Dashed lines are associated 95% pointwise confidence bands estimated as ± 2 standard errors.

If increases in volatility tend to reflect greater *idiosyncratic* rather than *systematic* variation, then volatility should only weakly predict returns. The finance literature has repeatedly documented a weak and inconsistent relation between volatility and expected returns (see Yu and Yuan, 2011, for a summary). Our finding is consistent with this literature.

Such patterns in the bottom decile of the return distribution could arise if we underestimate the hit probability when volatility is low. However, we are no more likely to underestimate hit probability when volatility is low than when it is high. For example, over the 105 months that the bottom-decile

DOI: 10.1002/jae

QPK exceeded 1.7, returns fell below our estimated conditional bottom decile in 10 of them, which is actually slightly less than 10%. Similarly, in the 98 months that the bottom-quartile QPK exceeded 1.35, returns fell below our estimated conditional bottom quartile in 25 of them (25.5%). Moreover, the moment condition implied by our quantile regression imposes zero correlation between hits and volatility.

We are unaware of any papers in the literature that report this negative correlation. Our findings are consistent with Figure 1 in Bakshi *et al.* (2010), which we reproduce in Figure A.2 in the online Appendix. Bakshi *et al.* use a nonparametric method to estimate the average projected pricing kernel in two separate periods. Their plots show a flatter pricing kernel in the high-volatility (1997–2007) period than in the low-volatility (1988–1996) period. Because their focus is on non-monotonicity, they do not discuss this result in the text of their paper. Song and Xiu (2016) obtain results consistent with ours too. They project the pricing kernel on both returns and volatility and find high estimates of the pricing kernel for low volatility.[15]

Our findings contradict recent models that specify the pricing kernel as a quadratic function of returns as in Christoffersen *et al.* (2013). In those models, state prices in both tails of the return distribution increase with volatility. Figure 7(f) shows that the projected pricing kernel tends to rotate anticlockwise when volatility increases.

6. CONNECTING SPOCQ TO RISK FACTORS

The preceding sections show that high return volatility is associated with lower state prices for negative returns and higher state prices for returns in the top quartile of the distribution. We interpret the changes in volatility for our sample as due more to changes in idiosyncratic than systematic volatility. If our interpretation is correct, then controlling for other variables that reflect the underlying state of the macroeconomy and financial markets that are related to the pricing kernel should not attenuate this correlation. In fact, we expect the relationship between the left-tail QPK and volatility to become more strongly negative after controlling for systematic risk factors. Thus we regress our monthly SPOCQ series on variables that capture (i) volatility, (ii) sentiment, (iii) liquidity and (iv) macroeconomic/financial risk.

6.1. Preliminaries

Expected returns increase when the projected pricing kernel rotates clockwise, so we expect variables that increase expected returns to have positive effects on the left-tail SPOCQ and negative effects on the right-tail SPOCQ conditional on the volatility of returns based on the expressions in equations (11) and (13). We base our set of regressors on those variables that have been shown to exhibit significant predictive power for expected excess stock returns.[16] Next, we briefly describe the variables we consider. Online Appendix A.6 provides further details.

(i) *Volatility & Tail Jumps*. In addition to the continuous-variation volatility variable we use in estimating SPOCQ, we also include a tail jump variable (Bollerslev and Todorov, 2011). We construct a measure of tail jumps using the log difference between the realized standard deviation and the realized continuous variation of the S&P 500 daily log returns over the previous 20 trading days.
(ii) *Sentiment*. Investor sentiment is a possibly erroneous belief about the future that may affect asset prices. Following Lemmon and Portniaguina (2006), we use the Michigan Consumer Sentiment

[15] See the bottom right panel of their Figure 10.
[16] We refer to the recent survey by Rapach and Zhou (2013), among others, for a compact discussion of the topic—admittedly, one of the most widely studied topics in empirical asset pricing.

DOI: 10.1002/jae

Index as our first sentiment measure. We also use the S&P 500 return over the past 20 trading days as a proxy for sentiment, based on the notion that changes in the pricing kernel in response to recent stock returns may reflect mood more than fundamentals.[17]

(iii) *Liquidity*. Evidence that less liquid assets have higher expected returns has been presented by numerous authors (e.g. Pástor and Stambaugh, 2003). We employ two measures of liquidity from our S&P 500 options data: one based on the logarithm of the dollar trading volume and one based on the logarithm of open interest on the relevant trading date. Our hypothesis is that, when the market is less liquid for options with strikes in a particular part of the distribution, then writers of those options require a greater premium. For example, a less liquid market for out-of-the-money call options may raise the price of those options, thereby raising the SPOCQ in the right tail of the return distribution. This is one potential explanation for the pricing kernel puzzle, although it would be incomplete without a theory for liquidity fluctuations. Alternatively, option buyers may require a premium to hold illiquid options because they fear an illiquid market if they try to sell the options before maturity. This alternative hypothesis implies lower SPOCQ for regions of the distribution with illiquid options.[18]

(iv) *Macroeconomic/Financial Risk*. We use five variables to capture macroeconomic or financial risk factors. First, we use the log of the dividend yield (Campbell and Shiller, 1988). As with the liquidity variables, an association between the dividend yield and the SPOCQ provides an incomplete explanation without a theory for what causes changes in dividend yield. Following Lettau and Ludvigson (2001), we use the variable known as CAY, which is the residual of the cointegrating relation for log consumption (C), log asset wealth (A) and log labor income (Y). Uncertainty about the future state of the economy may increase risk aversion, which is a prediction of the model in Pástor and Veronesi (2012). To account for this possibility, we include the daily news-based Economic Policy Uncertainty index from the Policy Uncertainty Project at Stanford University—we sum over the most recent 28 days and take the logarithm. Finally, we use two interest rate variables to capture macroeconomic conditions. We use the difference between BAA and AAA corporate bond yields as a measure of the default spread (Ferson and Harvey, 1991), and we use the difference between the 10-year and the 3-month Treasury yields to measure the term spread.

We use the QPK rather than SPOCQ as the dependent variable to make the coefficients comparable across models. We estimate our models using GLS as described in Section 3.3, and we standardize all explanatory variables to have mean zero and standard deviation one to ease comparison of their coefficients. To keep the discussion tractable and to avoid overfitting the data, we report results for three specifications. Specification 1 reproduces our volatility correlations by including volatility as the only regressor. In specification 2, we examine whether our liquidity, sentiment and macroeconomic/financial variables explain the QPK time variation as well as volatility. In specification 3, we check whether the regressors of specification 2 maintain their explanatory power once we add the two volatility variables. We refer to these three specifications, as SPEC1, SPEC2 and SPEC3, respectively.

6.2. Left-Tail Results

Table II presents results for the bottom decile (0–10%), bottom quartile (0–25%), and bottom half (0–50%) of the distribution. SPEC1 shows negative and statistically significant coefficients on volatility in all three cases, which restates the discussion in Section 5. SPEC2 shows that the sentiment,

[17] Other possibilities for sentiment include the closed-end fund discount (CEFD), the bull–bear spread based on the survey of American Association of Individual Investors, the Consumer Board Consumer Confidence Index, the Yale/Shiller crash confidence index, the Duke/CFO survey and the investor sentiment index in Baker and Wurgler. Preliminary regressions with these variables did not change our results qualitatively.

[18] We thank an anonymous reviewer for this suggestion.

DOI: 10.1002/jae

Table II. Quantile pricing kernel regressions: lower part of the return distribution

	0-10%			0-25%			0-50%		
Variable	SPEC1	SPEC2	SPEC3	SPEC1	SPEC2	SPEC3	SPEC1	SPEC2	SPEC3
Log CV	−0.093***		−0.156***	−0.049**		−0.088***	−0.063***		−0.080***
	(0.032)		(0.040)	(0.024)		(0.029)	(0.019)		(0.022)
Log tail jumps			0.027			0.001			(0.006)
			(0.020)			(0.012)			(0.006)
Returns		−0.087***	−0.085***		−0.001	−0.003		0.027***	0.018**
		(0.021)	(0.020)		(0.012)	(0.012)		(0.007)	(0.007)
Consumer sentiment		0.042	0.023		0.019	0.013		0.011	0.014
		(0.049)	(0.048)		(0.029)	(0.028)		(0.017)	(0.015)
Log open interest		−0.046	−0.076**		−0.068***	−0.076***		−0.038***	−0.039***
		(0.039)	(0.038)		(0.023)	(0.023)		(0.013)	(0.012)
Log dollar volume		0.125***	0.124***		0.069***	0.076***		0.031**	0.033***
		(0.038)	(0.036)		(0.022)	(0.021)		(0.012)	(0.012)
Log dividend yield		0.068*	−0.001		0.033	0.011		0.037***	0.027**
		(0.040)	(0.042)		(0.024)	(0.024)		(0.014)	(0.013)
CAY		0.036	0.046		−0.007	0.003		−0.017	−0.01
		(0.037)	(0.037)		(0.022)	(0.022)		(0.013)	(0.012)
Log uncertainty index sum		0.018	0.042		0.009	0.019		0.005	0.015
		(0.036)	(0.036)		(0.022)	(0.021)		(0.012)	(0.012)
Default spread		−0.007	0.069**		0.008	0.039**		−0.001	0.018*
		(0.027)	(0.032)		(0.017)	(0.019)		(0.010)	(0.010)
Term spread		0.060**	0.047**		0.064***	0.055***		0.038***	0.032***
		(0.024)	(0.024)		(0.014)	(0.014)		(0.008)	(0.007)
Constant	1.573***	1.572***	1.563***	1.246***	1.250***	1.243***	1.207***	1.230***	1.215***
	(0.031)	(0.029)	(0.030)	(0.026)	(0.025)	(0.026)	(0.026)	(0.023)	(0.025)
obs.	268	264	264	268	264	264	268	264	264
R-sq adj.	0.089	0.315	0.459	0.157	0.267	0.455	0.233	0.320	0.453
Wald test	7.406	141.437	268.674	14.042	93.017	301.756	27.152	120.110	377.967
	[0.006]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]

Note: We report GLS estimates and model diagnostics for the quantile pricing kernel regressions in Section 6 that pertain to the lower part of the return distribution. Standard errors are in parentheses and p-values are in squared brackets. The asterisks indicate statistical significance as follows: 1% (***), 5% (**), 10% (*). We standardize all explanatory variables such that their mean is zero and their standard deviation is one to ease comparison of their coeficients. SPEC2 and SPEC3 are based on a smaller number of observations due to lack of CAY data for the four months in 2012.

DOI: 10.1002/jae

liquidity and macro/finance variables have substantial explanatory power for the QPK. The adjusted R^2 is larger for SPEC2 than for SPEC1 in all three cases.

The coefficient on volatility is more negative in SPEC3 than SPEC1 for each of the three cases considered. This result suggests that the additional variables soak up some of the systematic component of volatility, so the volatility coefficient more strongly reflects idiosyncratic variation in volatility. Moreover, because we use trailing realized volatility rather than expected future volatility, and because volatility is mean reverting, the relationship between left-tail QPK and expected volatility would be even stronger than that depicted in Table II.[19] Together, these results are consistent with our interpretation that a negative volatility coefficient implies that much of the variation in volatility is idiosyncratic.

The lagged return, which we classify as a sentiment indicator, has a significant negative coefficient in the bottom decile. This result suggests that a high return over the past month reduces aversion to events in the left tail of the return distribution. However, the sentiment variables have small and mostly insignificant coefficients otherwise in Table II, so their relevance is mostly in the far left tail.

Our liquidity variables measure characteristics of trading in options with strikes in the relevant quantile, i.e. they differ across quantiles.[20] These variables enter the model with opposite signs and are statistically significant in SPEC3. The opposite signs and comparable magnitudes of these coefficients, along with the fact that these variables enter in logs, suggest that volume relative to open interest is what affects the SPOCQ. In particular, high volume relative to open interest is associated with greater aversion to downside risk. This finding suggests that out-of-the-money put option buyers demand a liquidity premium, or that heavy trading activity in the tails is symptomatic of higher risk aversion.

SPEC2 produces results for the dividend yield that appear consistent with the return predictability literature (e.g. Campbell and Shiller, 1988). A high dividend yield indicates a high value of left-tail states and therefore high expected returns. For both the lowest decile and quartile, this coefficient is insignificant once we control for volatility in SPEC3. It remains significant in the 0–50% model. The CAY and policy uncertainty variables are statistically insignificant in all cases.[21]

The term spread is positively associated with the left-tail pricing kernel (SPEC2 and SPEC3), which implies that steepening of the term structure slope coincides with high risk aversion. SPEC3 shows that a high default spread implies high state prices in all three cases, which is consistent with the findings of Engle and Rosenberg (2002) for countercyclicality of risk aversion. However, this result depends on volatility being in the model. When we remove volatility (SPEC2), the coefficients on the default spread switch sign, suggesting that the explanatory power of the default spread is easily confounded with that of volatility for the bottom decile and bottom half of the distribution.

6.3. Right-Tail Results and the Pricing Kernel Puzzle

Table III contains the right-tail counterparts of the left-tail regressions in Table II. For SPEC1, we see the positive correlation between volatility and the right-tail SPOCQ, especially in the top quartile. The change in the volatility coefficient is small when we control for the tail jump, sentiment, liquidity and macro/finance variables. Based on the results for SPEC2, the sentiment, liquidity and macro/finance variables have substantial explanatory power for the right-tail QPK. The adjusted R^2 for SPEC2 is much larger than for SPEC1 for the top-decile regressions and similar to SPEC1 for the top-quartile regression.

[19] We thank an anonymous reviewer for raising this point.

[20] As we discuss in detail in online Appendix A.6, for each trading date we calculated the moneyness for each put and call strike for options with 1 month to expiration. We then assigned the open interest for the call and put options in the six moneyness bins implied by the estimates of the conditional quantiles in Section 4.

[21] The CAY variable is measured during the month whereas the QPK is measured on a particular day mid month. This difference could explain some of the weakness in explanatory power. A similar point could be made for the consumer sentiment index.

DOI: 10.1002/jae

Table III. Quantile pricing kernel regressions: upper part of the return distribution

	75-100%			90-100%		
Variable	SPEC1	SPEC2	SPEC3	SPEC1	SPEC2	SPEC3
Log CV	0.101***		0.108***	0.061		0.050
	(0.033)		(0.036)	(0.051)		(0.059)
Log tail jumps			0.013			−0.009
			(0.009)			(0.015)
Returns		−0.016	−0.011		-0.017	−0.016
		(0.010)	(0.010)		(0.016)	(0.016)
Consumer sentiment		−0.044*	−0.050**		−0.067**	−0.067**
		(0.025)	(0.021)		(0.033)	(0.033)
Log open interest		0.008	0.007		0.045*	0.046*
		(0.013)	(0.013)		(0.027)	(0.027)
Log dollar volume		0.044***	0.034**		0.028	0.024
		(0.016)	(0.015)		(0.025)	(0.026)
Log dividend yield		−0.072***	−0.068***		−0.098***	−0.094***
		(0.022)	(0.018)		(0.028)	(0.028)
CAY		0.057***	0.053***		0.133***	0.129***
		(0.018)	(0.015)		(0.024)	(0.025)
Log uncertainty index sum		−0.036*	−0.045***		−0.047*	−0.050*
		(0.019)	(0.016)		(0.026)	(0.027)
Default spread		0.010	−0.008		0.012	0.002
		(0.016)	(0.014)		(0.020)	(0.023)
Term spread		−0.036***	−0.029***		−0.021	−0.018
		(0.014)	(0.011)		(0.016)	(0.017)
Constant	0.636***	0.674***	0.712***	0.509***	0.577***	0.581***
	(0.044)	(0.039)	(0.043)	(0.060)	(0.063)	(0.064)
Obs.	268	264	264	268	264	264
R^2	0.495	0.464	0.663	−0.001	0.242	0.236
Wald test	63.490	215.787	915.461	0.216	106.213	134.238
	[0.000]	[0.000]	[0.000]	[0.642]	[0.000]	[0.000]

Note: We report GLS estimates and model diagnostics for the quantile pricing kernel regressions in Section 6 that pertain to the upper part of the return distribution. Standard errors are in parentheses and p-values are in squared brackets. The asterisks indicate statistical significance as follows: 1% (***), 5% (**), 10% (*). We standardize all explanatory variables such that their mean is zero and their standard deviation is one to ease comparison of their coeficients. SPEC2 and SPEC3 are based on a smaller number of observations due to lack of CAY data for the four months in 2012.

Unlike the left tail, the top-decile quantile estimates in Table I failed some diagnostic tests, which means that some variation in the top-decile SPOCQ could stem from changes in the probability of right-tail events rather than changes in the projected pricing kernel. For this reason, we focus the discussion on the top-quartile results.

The tail-jump variable is unrelated to the right-tail QPK, and consumer sentiment exhibits a negative relationship with the right-tail pricing kernel. Regarding the liquidity variables, volume enters with a significant positive coefficient for the top quartile, but open interest generally has a small and insignificant coefficient. Thus high volume is associated with greater aversion to right-tail risk. As with the left tail, these results suggest that out-of-the-money call option buyers demand a liquidity premium, or that heavy trading activity in the tails is symptomatic of higher state prices.

Expected returns increase when the projected pricing kernel rotates clockwise, placing a larger discount on negative returns and a smaller discount on positive returns. Consistent with the return predictability literature, increases in the dividend yield or term spread are associated with a clockwise rotation in the pricing kernel, thereby increasing expected returns conditional on the volatility of returns. These variables enter with negative coefficients in the right tail and positive coefficients in the left tail. The CAY variable has a significant positive coefficient in the right tail and a positive

DOI: 10.1002/jae

but insignificant coefficient in the bottom decile. Thus our results do not suggest a clear association between CAY and expected returns.

7. CONCLUSION

We track the option-implied stochastic discount factor for S&P 500 returns between January 1990 and April 2012. We do so by developing a set of statistics, the state prices of conditional quantiles (SPOCQ), which reflect the market's willingness to pay for insurance against outcomes across various quantiles of the return distribution. We construct the SPOCQ series by evaluating the risk-neutral return distribution at the conditional quantiles of the physical return distribution. There is a direct and straightforward link between the pricing kernel and SPOCQ. The projected pricing kernel over a range of the return distribution is equal to the SPOCQ divided by the probability mass associated with such a range. We label this ratio the quantile pricing kernel (QPK).

To our knowledge, we offer the first attempt in the literature to examine the time variation in the discount factor using quantiles. Consistent with the robust evidence presented in other studies, we find that the projected pricing kernel is non-monotonic. Moreover, our focus on quantiles allows us to discover that volatility relates to the projected pricing kernel in a surprising way. High volatility is associated with a relatively low value of left-tail returns. This finding implies that much of the variation in volatility for our sample was idiosyncratic to the S&P 500 and not associated with the pricing kernel.

ACKNOWLEDGEMENTS

We thank two anonymous referees for comments that significantly improved the paper. Any remaining errors are ours.

REFERENCES

Azzalini A, Capitanio A. 2003. Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t distribution. *Journal of the Royal Statistical Society, Series B* **65**: 367–389.

Bakshi G, Madan D, Panayotov G. 2010. Returns of claims on the upside and the viability of U-shaped pricing kernels. *Journal of Financial Economics* **97**: 130–154.

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

Breeden D, Litzenberger R. 1978. Prices of state-contingent claims implicit in option prices. *Journal of Business* **51**: 621–651.

Buchinsky M. 1998. Recent advances in quantile regression models: a practical guideline for empirical research. *Journal of Human Resources* **33**: 88–126.

Campbell J, Shiller R. 1988. The dividend–price ratio and expectations of future dividends and discount factors. *Review of Financial Studies* **1**: 195–228.

Chamberlain G. 1982. Multi-variate regression models for panel data. *Journal of Econometrics* **18**: 5–46.

Cho J, White H. 2007. Testing for regime switching. *Econometrica* **75**: 1671–1720.

Christoffersen P, Heston S, Jacobs K. 2013. Capturing option anomalies with a variance-dependent pricing kernel. *Review of Financial Studies* **26**: 1962–2006.

Engle R, Manganelli S. 2004. CAViaR: conditional autoregressive value at risk by regression quantiles. *Journal of Business and Economic Statistics* **22**: 367–381.

Engle R, Rosenberg J. 2002. Empirical pricing kernels. *Journal of Financial Economics* **64**: 341–372.

Ferson W, Harvey C. 1991. The variation of economic risk premiums. *Journal of Political Economy* **99**: 385–415.

Figlewski S. 2010. Estimating the implied risk neutral density for the U.S. market portfolio. In *Volatility and Time Series Econometrics: Essays in Honor of* , Robert F Engle, Bollerslev T, Russell J, Watson M (eds), Ch. 15. Oxford University Press: Oxford; 323–353.

French K, Schwert G, Stambaugh R. 1987. Expected returns and volatility. *Journal of Financial Economics* **100**: 3–299.

Heston S, Nandi S. 2000. A closed-form GARCH option pricing model. *Review of Financial Studies* **13**: 585–626.

DOI: 10.1002/jae

Jackwerth J. 2000. Recovering risk aversion from option prices and realized returns. *Review of Financial Studies* **13**: 433–451.
Koenker R, Bassett G. 1978. Regression quantiles. *Econometrica* **46**: 33–50.
Lemmon M, Portniaguina E. 2006. Consumer confidence and asset prices: some empirical evidence. *Review of Financial Studies* **19**: 1499–1529.
Lettau M, Ludvigson S. 2001. Consumption, aggregate wealth, and expected stock returns. *Journal of Finance* **56**: 815–849.
Marron J, Wand M. 1992. Exact mean integrated squared error. *Annals of Statistics* **20**: 712–736.
Nevo A. 2001. Measuring market power in the ready-to-eat cereal industry. *Econometrica* **69**: 307–342.
Pástor L, Stambaugh R. 2003. Liquidity risk and expected stock returns. *Journal of Political Economy* **111**: 642–685.
Pástor L, Veronesi P. 2012. Uncertainty about government policy and stock prices. *Journal of Finance* **67**: 1219–1264.
Rapach D, Zhou G. 2013. Forecasting stock returns. In *Handbook of Economic Forecasting*, Elliott G, Timmermann A (eds), Vol. 2, Part A. North-Holland: Amsterdam; 328–383.
Song Z, Xiu D. 2016. A tale of two option markets: state-price densities implied from S&P 500 and VIX option prices. *Journal of Econometrics* **190**: 176–196.
White H, Kim T, Manganelli S. 2010. Modeling autoregressive conditional skewness and kurtosis with multi-quantile CAViaR. In *Volatility and Time Series Econometrics: Essays in Honor of*, Engle F, Bollerslev T, Russell J, Watson M (eds). Oxford University Press: Oxford; 231–256.
Yu J, Yuan Y. 2011. Investor sentiment and the mean–variance relation. *Journal of Financial Economics* **100**: 367–381.

DOI: 10.1002/jae