The Term Structures of Expected Loss and Gain Uncertainty

Bruno Feunou*
Bank of Canada

Ricardo Lopez Aliouchkin
Syracuse University

Roméo Tédongap
ESSEC Business School

Lai Xu
Syracuse University

November 2019

Abstract

We document that the term structures of risk-neutral expected loss and gain uncertainty on the S&P500 returns are upward sloping on average. These shapes mainly reflect the higher premium required by investors to hedge downside risk, and the belief that potential gains will increase in the long-run. The term structures exhibit substantial time series variation with large negative slopes during crisis periods. Through the lens of Andersen et al. (2015)'s framework, we evaluate the ability of existing reduced-form option pricing models to replicate these term structures. We stress that three ingredients are particularly important: (1) the inclusion of jumps; (2) disentangling the price of negative jump risk from its positive analog in the stochastic discount factor specification; (3) specifying three latent factors.

Keywords: Quadratic payoff, quadratic loss, quadratic gain, quadratic risk premium, options

JEL Classification: G12

*Corresponding Author: Bank of Canada, 234 Wellington St., Ottawa, Ontario, Canada K1A 0G9. Tel: +1 613 782 8302. Email: feun@bankofcanada.ca. Feunou gratefully acknowledges financial support from the IFSID. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of Canada.

1 Introduction

Financial economists have long agree that in order to better understand asset returns, but also uncertainty about these returns, it would be necessary to break them down into several components, each reflecting a different aspect through which an investment opportunity can be perceived, analyzed and evaluated. Since a return (r) can be classified as either a loss ($-l$) if nonpositive or a gain (g) if nonnegative, it is then natural to break it down into these two components, formally $r = g - l$ where $l = \max(0, -r)$ and $g = \max(0, r)$. This decomposition of returns leads also to a similar decomposition of return uncertainty into loss and gain components, namely loss uncertainty and gain uncertainty (also referred to as downside and upside, respectively, or in the most recent literature as bad and good, respectively. See for example Barndorff-Nielsen et al., 2010, Patton and Sheppard, 2015, Bekaert et al., 2015, and Kilic and Shaliastovich, 2019, just to name a few). Likewise, investment returns are assessed over a given horizon which together with the maturity of the payoff are part of the key elements that guide investment choices.

We argue that expectations of (per-period) asset returns uncertainty and its loss and gain components across different investment horizons, i.e. their respective term structures, are critical for several financial decisions and risk management. It is also important when dealing with expectations of asset returns uncertainty to distinguish between physical expectations and risk-neutral expectations. On one hand, physical expectations of uncertainty measure anticipations of how much investors believe they could be wrong about their returns forecast. On the other hand, risk-neutral expectations of uncertainty additionally inform about how much investors are willing to pay for risk hedging or would be willing to require for risk compensation.

The primary goal of this article is an empirical investigation of physical and risk-neutral expectations of loss uncertainty and gain uncertainty across different investment horizons. The main challenge resides in estimating or measuring these expectations using available financial data. Since current period uncertainty on a future period return is not observed, a large body of the literature relies on model-free measures that can readily be computed using realized returns. A popular measure of uncertainty is the realized variance that cumulates higher frequency squared returns over the investment horizon. The loss component of realized variance cumulates higher frequency squared losses while the gain component sums up higher frequency squared gains over the investment horizon. Thus, the realized variance is the sum of its loss and gain components. As thoroughly discussed in Feunou et al. (2019), estimating or measuring risk-neutral expectations of loss and gain realized variance is neither feasible, nor are loss and gain variance swaps traded such that their strikes could then be observed measures of these risk-neural expectations. It is therefore important to rely on a measure of asset returns uncertainty for which the unobserved expectations of loss and gain components can be consistently estimated or measured from the data.

Unlike the realized variance, the quadratic payoff is the squared of the realized return over the investment horizon. It is also a measure of the asset returns uncertainty. The loss component of the quadratic payoff (or quadratic loss) is the squared loss while the gain component (or quadratic gain) is the squared gain over the investment horizon. Similar to realized variance, the quadratic payoff is the sum of its loss and gain components. To the contrary of realized variance, both physical and risk-neutral expectations of quadratic loss and gain can be consistently estimated or measured from the data. We provide more details in the Internet Appendix Section A.1. We therefore rely on the quadratic payoff when analyzing the term structure of expected loss uncertainty and gain uncertainty.

Using a large panel of S&P500 index options data with time-to-maturity ranging from one month to twelve months, we build model-free risk-neutral expected quadratic loss and gain term structures. Our methodology follows from Bakshi et al. (2003) and is similar to that used to compute the VIX index. Likewise, using high-frequency S&P500 index return data and relying on a state-of-the-art variance forecasting model considered by Bekaert and Hoerova (2014), we build physical expected quadratic loss and gain term structures. We ask to what extent variations in these term structures reflect changes in the anticipated path of future loss and gain uncertainty, and therefore, the extent to which they reflect changes in the risk premia.

Our results reveal new important findings. First, the average term structure of the physical expected quadratic loss is downward sloping (a slope of -4.73 percent-square), while the average term structure of the risk-neutral expected quadratic loss is upward sloping (a slope of 3.63 percent-square). This means that, on average, investors anticipate that the (per-period) loss potential decreases with the investment horizon, yet at the same time on the market, hedging the long-term loss potential is more expensive than hedging the short-term loss potential of stocks. Second, the average term structure of the physical expected quadratic gain is upward sloping (a slope of 7.09 percent-square), while the average term structure of the risk-neutral expected quadratic gain is slightly upward sloping, almost flat (a slope of 1.01 percent-square). Likewise, this means that, on average, investors foresee that the (per-period) gain potential increases with the investment horizon, yet at the same time on the market, speculating on the short-term gain potential is almost as costly as speculating on the long-term gain potential of stocks.

Our estimates of physical and risk-neutral expectations of quadratic loss and quadratic gain allow us to compute the associated risk premia by taking the appropriate difference between the physical and the risk-neutral expectation. We follow Feunou et al. (2019) and measure the loss quadratic risk premium (QRP) as the risk-neutral minus the physical expected quadratic loss. It is the premium paid for downside risk hedging and thus a measure of downside risk. Likewise, we measure gain QRP as the physical minus the risk-neutral expected quadratic gain. It is the premium received for upside risk compensation and thus a measure of upside risk. We subsequently analyze the term structures of loss and gain QRPs, and we obtain that both term structures are upward sloping (slopes of 8.36 percent-square and 6.09 percent-square, respectively). Therefore, on average, the (per-period) downside and upside risks are both higher for long-term investments relative to short-term investments in stocks, and since the equity premium is a remuneration of both types of risk, this confirms the upward sloping average term structure of the equity premium found elsewhere in the literature.

The secondary goal of this article is to evaluate whether leading option pricing models which predominantly appear to be special cases of the model of Andersen et al. (2015) (henceforth AFT) are able to replicate the actual term structures of risk-neutral expected quadratic loss and gain. Key features of the three-factor AFT model are its flexibility and its ability to completely disentangle the negative from the positive jump dynamics. To enhance our understanding of the model ingredients underlying the statistical properties of the quadratic loss and gain, we also estimate several restricted variants of the AFT model. These include, among others, the two-factor diffusion model of Christoffersen et al. (2009) (denoted as the baseline model AFT0), and a version of the AFT model where the negative and positive jump dynamics are equal (denoted by AFT3). The AFT3 model essentially represents the vast majority of option and variance swaps models studied in the literature so far (see e.g., Bates, 2012, Christoffersen et al., 2012, Eraker, 2004, Chernov et al., 2003, Huang and Wu, 2004, Amengual and Xiu, 2018, and Ait-Sahalia et al., 2015).[1] We find that accounting for jumps in asset prices is essential for the model to fit the term structure of the risk-neutral expected quadratic loss and gain. The AFT0 model overestimates the risk-neutral expected quadratic gain and underestimates the risk-neutral expected quadratic loss, but is able

[1]Some authors consider asymmetry in the jump size distribution (see e.g., Amengual and Xiu, 2018) However, the jump size distribution is assumed to be constant and the time variation in jumps comes through the jump intensity which is assumed to be the same regardless of the sign of the jump.

to fit the term structure of the risk-neutral expected quadratic payoff. We also find that a jump process rather than a diffusion process is the most important in fitting the term structure of the risk-neutral expected quadratic loss, while it appears to be the opposite for the term structure of the risk-neutral expected quadratic gain.

The AFT model is primarily used for the risk-neutral dynamics of asset prices, and we further couple it with a pricing kernel specification that maps the risk-neutral into the physical dynamics. All parameters are estimated to maximize the joint likelihood of risk-neutral expected quadratic loss and gain across the term structure together with the second and third risk-neutral cumulants of asset returns. We examine the ability of various pricing kernel specifications in matching the actual dynamics of the term structures of physical expected quadratic loss and gain. This is equivalent to matching the actual term structures of loss and gain QRP. Our results unequivocally point to the importance of disentangling the price of negative jumps from the price of positive jumps. In other words, a restricted version of the pricing kernel imposing the same price for the negative and positive jump risk is unable to match the dynamics of the loss and gain QRP together. This restricted version represents the vast majority of pricing kernels studied in the literature (see for instance, Eraker, 2004, Santa-Clara and Yan, 2010, Christoffersen et al., 2012 and Bates, 2012), and highlights its inability to account for the joint actual dynamics of the loss and gain QRP.

Our paper is related to the recent literature that analyzes the term structure of variance swaps. Ait-Sahalia et al. (2015) and Amengual and Xiu (2018) specify reduced-form models for the term structure of total variance. Dew-Becker et al. (2017) investigates the ability of existing structural models in fitting the observed term structure of variance swaps. We contribute to this literature by investigating the term structures of the two variance components. Our paper also relates to another strand of the literature documenting the importance of analyzing loss and gain components of variance (risk-neutral or physical) and VRP. Barndorff-Nielsen et al. (2010) provide the theoretical arguments supporting the splitting of the total realized variance into a loss and gain components.

The remainder of the paper is organized as follows. Section 2 introduces definitions and notations of all quantities, the data, and the methodology for constructing the risk-neutral and physical term structure of expected quadratic loss and gain, and presents key empirical facts that any economically sound model should be able to replicate. Section 3 introduces the AFT model and provides some details on its properties, including the implied closed-form for both the risk-neutral and physical expectation of the quadratic loss and gain. Section 4 provides detail on the estimation on the AFT model. Section 5 evaluates the ability of the AFT model and its variants in fitting the empirical facts. Section 6 concludes.

2 Methodology, Data and Preliminary Analysis

In this section we start by introducing the quadratic payoff, and its loss and gain components, namely, the quadratic loss and the quadratic gain. Next, we introduce a heuristic theoretical framework to understand the difference between the quadratic loss and the quadratic gain. We discuss the methodology to measure the risk-neutral and physical expectations of the quadratic payoff, the quadratic loss and the quadratic gain, over a given investment horizon. For the purpose of computing these term structures, we present the data and provide descriptive statistics. Finally, we provide a preliminary analysis based on principal components extracted from these term structures.

2.1 Definitions

Let S_t denote the S&P 500 index price at the end of day t, and for any investment horizon τ, let $r_{t,t+\tau}$ denote its (log) return from end of day t to end of day $t+\tau$, given by $r_{t,t+\tau} = \ln\left(S_{t+\tau}/S_t\right)$. Both the log return $r_{t,t+\tau}$ and the quadratic payoff $r^2_{t,t+\tau}$ are subject to a gain-loss decomposition as follows:

$$r_{t,t+\tau} = g_{t,t+\tau} - l_{t,t+\tau} \quad \text{and} \quad r^2_{t,t+\tau} = g^2_{t,t+\tau} + l^2_{t,t+\tau}, \tag{1}$$

where the gain $g_{t,t+\tau} = \max\left(0, r_{t,t+\tau}\right)$ and the loss $l_{t,t+\tau} = \max\left(0, -r_{t,t+\tau}\right)$, represent the positive and negative parts of the asset payoff, respectively. In other words, the gain and the loss are nonnegative amounts flowing in and out of an average investor's wealth, respectively. Since a positive gain and a positive loss cannot occur simultaneously, we observe that $g_{t,t+\tau} \cdot l_{t,t+\tau} = 0$. This gain loss decomposition of an asset's payoff is exploited as an asset pricing context by Bernardo and Ledoit (2000).

Our goal in this article is to study how the time series dynamics of risk-neutral expectations $\mathbb{E}^{\mathbb{Q}}_t\left[l^2_{t,t+\tau}\right]$ and $\mathbb{E}^{\mathbb{Q}}_t\left[g^2_{t,t+\tau}\right]$, and of the physical expectations $\mathbb{E}^{\mathbb{P}}_t\left[l^2_{t,t+\tau}\right]$ and $\mathbb{E}^{\mathbb{P}}_t\left[g^2_{t,t+\tau}\right]$, vary with the investment horizon τ, where the exponents $^{\mathbb{Q}}$ and $^{\mathbb{P}}$ indicate that the values are under the risk-neutral and the physical measures, respectively. Knowledge of these term structures can be relevant in various risk management contexts. Indeed, one can learn about investors' anticipations of the degree of loss and gain uncertainty every day for each investment horizon, and also how much investors are willing to pay for hedging, or to require for compensation of the associated risks over a given investment horizon.

Given the risk-neutral and physical expectations of the same random quantity, one can readily take their difference to measure the associated risk premium. Following Feunou et al. (2019), we define the difference between the risk-neutral and the physical expectations of the quadratic payoff as the quadratic risk premium (QRP), which the loss and gain components, called loss QRP and gain QRP, and denoted by $QRP_t^l(\tau)$ and $QRP_t^g(\tau)$, respectively, are formally given by:

$$QRP_t^l(\tau) \equiv \mathbb{E}_t^{\mathbb{Q}}\left[l_{t,t+\tau}^2\right] - \mathbb{E}_t^{\mathbb{P}}\left[l_{t,t+\tau}^2\right] \quad \text{and} \quad QRP_t^g(\tau) \equiv \mathbb{E}_t^{\mathbb{P}}\left[g_{t,t+\tau}^2\right] - \mathbb{E}_t^{\mathbb{Q}}\left[g_{t,t+\tau}^2\right]. \tag{2}$$

Equation (2) shows that loss QRP (QRP^l) represents the premium paid for the insurance against fluctuations in loss uncertainty, while the gain QRP (QRP^g) is the premium earned to compensate for the fluctuations in gain uncertainty. Thus, the (net) QRP ($QRP \equiv QRP^l - QRP^g$) represents the net cost of insuring fluctuations in loss uncertainty, that is the premium paid for the insurance against fluctuations in loss uncertainty net of the premium earned to compensate for the fluctuations in gain uncertainty. Our study of the term structures of the risk-neutral and the physical expected quadratic loss and gain naturally leads to examining the term structures of loss and gain QRPs.

2.2 Dissecting the Quadratic Payoff into Loss and Gain: A Theory

For simplicity, let us denote the risk-neutral and physical expectations as the following:

$$\mu_n^{\mathbb{Q}+}(t,\tau) \equiv \mathbb{E}_t^{\mathbb{Q}}\left[g_{t,t+\tau}^n\right], \quad \mu_n^{\mathbb{Q}-}(t,\tau) \equiv \mathbb{E}_t^{\mathbb{Q}}\left[l_{t,t+\tau}^n\right], \quad \text{and } \mu_n^{\mathbb{Q}}(t,\tau) \equiv \mathbb{E}_t^{\mathbb{Q}}\left[r_{t,t+\tau}^n\right], \tag{3}$$

$$\mu_n^{\mathbb{P}+}(t,\tau) \equiv \mathbb{E}_t^{\mathbb{P}}\left[g_{t,t+\tau}^n\right], \quad \mu_n^{\mathbb{P}-}(t,\tau) \equiv \mathbb{E}_t^{\mathbb{P}}\left[l_{t,t+\tau}^n\right], \quad \text{and } \mu_n^{\mathbb{P}}(t,\tau) \equiv \mathbb{E}_t^{\mathbb{P}}\left[r_{t,t+\tau}^n\right], \tag{4}$$

To understand the difference between $\mu_2^{\mathbb{Q}+}(t,\tau)$ and $\mu_2^{\mathbb{Q}-}(t,\tau)$, we follow Duffie et al. (2000),

$$\begin{aligned} \mu_2^{\mathbb{Q}-}(t,\tau) &= \frac{\mathbb{E}_t^{\mathbb{Q}}\left[r_{t,t+\tau}^2\right] + \Lambda^{\mathbb{Q}}(t,\tau)}{2} \\ \mu_2^{\mathbb{Q}+}(t,\tau) &= \frac{\mathbb{E}_t^{\mathbb{Q}}\left[r_{t,t+\tau}^2\right] - \Lambda^{\mathbb{Q}}(t,\tau)}{2}, \end{aligned} \tag{5}$$

where $\Lambda^{\mathbb{Q}}(t,\tau)$, the wedge between the risk-neutral expected quadratic loss and gain, is given by:

$$\Lambda^{\mathbb{Q}}(t,\tau) = \frac{2}{\pi}\int_0^{+\infty} \frac{Im\left(\varphi_{t,\tau}^{(2)}(-iv)\right)}{v} dv \tag{6}$$

with $\varphi_{t,\tau}(\cdot)$ being the time t conditional risk-neutral moment-generating function of $r_{t,t+\tau}$, and $\varphi^{(2)}_{t,\tau}(\cdot)$ its second order derivative and $Im()$ refers to the imaginary coefficient of a complex number. From equation (5), it is apparent that, studying the term structure of $\mu_2^{\mathbb{Q}-}(t,\tau)$ and $\mu_2^{\mathbb{Q}+}(t,\tau)$ amounts to studying the term structure of the quadratic payoff $\mathbb{E}_t^{\mathbb{Q}}\left[r^2_{t,t+\tau}\right]$ and the term structure of $\Lambda^{\mathbb{Q}}(t,\tau)$. Several papers in the literature have already dealt successfully with $\mathbb{E}_t^{\mathbb{Q}}\left[r^2_{t,t+\tau}\right]$, and the consensus seems to be that a two-factor diffusion model provide a good statistical representation (see Christoffersen et al., 2009). We now try to understand conceptually the potential drivers of the wedge $\Lambda^{\mathbb{Q}}(t,\tau)$.

We use the following power series expansion of the moment generating function $\varphi_{t,\tau}(\cdot)$:

$$\varphi_{t,\tau}(v) = \sum_{n=0}^{\infty} \frac{v^n}{n!}\mu_n^{\mathbb{Q}}(t,\tau),$$

to establish that

$$\Lambda^{\mathbb{Q}}(t,\tau) = \lim_{\bar{v}\to\infty}\left\{\sum_{j=1}^{\infty}\frac{(-1)^j\,\bar{v}^{2j-1}}{(2j-1)(2j-1)!}\mu_{2j+1}^{\mathbb{Q}}(t,\tau)\right\}. \tag{7}$$

which is a weighted average of odd high order non-central moments. Since only the odd high moments are included, the wedge $\Lambda^{\mathbb{Q}}(t,\tau)$ is closely related to the asymmetry in the distribution of $r_{t,t+\tau}$. In the summation, when focusing on $j=1$, it is apparent that $\Lambda^{\mathbb{Q}}(t,\tau)$ is the opposite of the third order non-central moment $\mu_3^{\mathbb{Q}}(t,\tau)$ (up to a positive multiplicative constant). Recall that $\mu_3^{\mathbb{Q}}(t,\tau)$ is related to the first three central moments as following:

$$\mu_3^{\mathbb{Q}}(t,\tau) = \kappa_3^{\mathbb{Q}}(t,\tau) + 3\mu_1^{\mathbb{Q}}(t,\tau)\,\kappa_2^{\mathbb{Q}}(t,\tau) + \left[\mu_1^{\mathbb{Q}}(t,\tau)\right]^3,$$

were $\kappa_n^{\mathbb{Q}}(t,\tau) \equiv \mathbb{E}_t^{\mathbb{Q}}\left[\left(r_{t,t+\tau} - \mu_1^{\mathbb{Q}}(t,\tau)\right)^n\right]$. Hence, we conclude that the wedge between the risk-neutral expected quadratic loss and gain increases with the asymmetry in the risk-neutral distribution. A negative skewness implies larger risk-neutral expected quadratic losses, while a positive skewness yields the opposite effect. The wedge between the risk-neutral expected quadratic loss and gain still exist and always negative when the distribution is symmetric (all odd order central moments for a symmetric distribution are zero). In that case, the wedge increases in absolute value as the volatility increases.

2.3 Constructing Expectations

2.3.1 Inferring the Risk-neutral Expectation from Option Prices

In practice, previous literature estimates the risk-neutral conditional expectation of quadratic payoff directly from a cross-section of option prices. Bakshi et al. (2003) provide model-free formulas linking the risk-neutral moments of stock returns to explicit portfolios of options. These formulas are based on the basic notion, first presented in Bakshi and Madan (2000), that any payoff over a time horizon can be spanned by a set of options with different strikes with the same maturity as the investment horizon.

We adopt the notation in Bakshi et al. (2003), and define $V_t(\tau)$ as the time-t price of the τ-maturity quadratic payoff on the underlying stock. Bakshi et al. (2003) show that $V_t(\tau)$ can be recovered from the market prices of out-of-the-money (OTM) call and put options as follows:

$$V_t(\tau) = \int_{S_t}^{\infty} \frac{1 - \ln(K/S_t)}{K^2/2} C_t(\tau; K)\, dK + \int_{0}^{S_t} \frac{1 + \ln(S_t/K)}{K^2/2} P_t(\tau; K)\, dK, \tag{8}$$

where S_t is the time-t price of underlying stock, and $C_t(\tau; K)$ and $P_t(\tau; K)$ are time-t option prices with maturity τ and strike K, respectively. The risk-neutral expected quadratic payoff is then

$$\mathbb{E}_t^{\mathbb{Q}}\left[r_{t,t+\tau}^2\right] = e^{r_f \tau} V_t(\tau), \tag{9}$$

where r_f is the continuously compounded interest rate.

We compute $V_t(\tau)$ on each day and maturity. In theory, computing $V_t(\tau)$ requires a continuum of strike prices, while in practice we only observe a discrete and finite set of them. Following Jiang and Tian (2005) and others, we discretize the integrals in equation (8) by setting up a total of 1001 grid points in the moneyness (K/S_t) range from 1/3 to 3. First, we use cubic splines to interpolate the implied volatility inside the available moneyness range. Second, we extrapolate the implied volatility using the boundary values to fill the rest of the grid points. Third, we calculate option prices from these 1001 implied volatilities using the Black-Scholes formula proposed by Black and Scholes (1973).[2] Next, we compute $V_t(\tau)$ if there are four or more OTM option implied volatilities (e.g. Conrad et al., 2013 and others). Lastly, for example, to obtain $V_t(30)$ on a given day, we interpolate and extrapolate $V_t(\tau)$ with different τ. This process yields a daily time series of the risk-neutral expected quadratic payoff for each maturity $\tau = 30, 60, \ldots, 360$ days.

[2]Since the S&P 500 options are European, we do not have issues with the early exercise premium.

Note that the price of the quadratic payoff $V_t(\tau)$ in equation (8) is the sum of a portfolio of OTM call options and a portfolio of OTM put options:

$$V_t(\tau) = V_t^g(\tau) + V_t^l(\tau), \tag{10}$$

where:

$$V_t^l(\tau) = \int_0^{S_t} \frac{1+\ln(S_t/K)}{K^2/2} P_t(\tau; K)\, dK \quad \text{and} \quad V_t^g(\tau) = \int_{S_t}^{\infty} \frac{1-\ln(K/S_t)}{K^2/2} C_t(\tau; K)\, dK. \tag{11}$$

Feunou et al. (2019) analytically prove that $V_t^l(\tau)$ is the price of the quadratic loss, and $V_t^g(\tau)$ is the price of the quadratic gain. We present that proof in the Internet Appendix accompanying this paper. Hence, the risk-neutral expectation of quadratic loss and gain are:

$$\mathbb{E}_t^{\mathbb{Q}}\left[l_{t,t+\tau}^2\right] = e^{r_f\tau} V_t^l(\tau) \quad \text{and} \quad \mathbb{E}_t^{\mathbb{Q}}\left[g_{t,t+\tau}^2\right] = e^{r_f\tau} V_t^g(\tau). \tag{12}$$

2.3.2 Estimating the Physical Conditional Expected Quadratic Payoff

A regression model can be used to estimate the expectations of the quadratic payoff and truncated returns over different periods using actual returns data. To compute these expectations, we assume that, conditional on time-t information, log returns $r_{t,t+\tau}$ have a normal distribution with mean $\mu_{t,\tau} = \mathbb{E}_t[r_{t,t+\tau}] = Z_t^\top \beta_\mu$ and variance $\sigma_{t,\tau}^2 = \mathbb{E}_t[RV_{t,t+\tau}]$, where $\mathbb{E}_t[RV_{t,t+\tau}] = Z_t^\top \beta_\sigma$ and $RV_{t,t+\tau}$ is the realized variance between end of day t and end of day $t+\tau$. We have then:

$$\mathbb{E}_t\left[r_{t,t+\tau}^2\right] = \mu_{t,\tau}^2 + \sigma_{t,\tau}^2 \quad \text{and} \quad \begin{cases} \mathbb{E}_t\left[l_{t,t+\tau}^2\right] &= \left(\mu_{t,\tau}^2 + \sigma_{t,\tau}^2\right)\Phi\left(-\frac{\mu_{t,\tau}}{\sigma_{t,\tau}}\right) - \mu_{t,\tau}\sigma_{t,\tau}\phi\left(\frac{\mu_{t,\tau}}{\sigma_{t,\tau}}\right) \\ \mathbb{E}_t\left[g_{t,t+\tau}^2\right] &= \left(\mu_{t,\tau}^2 + \sigma_{t,\tau}^2\right)\Phi\left(\frac{\mu_{t,\tau}}{\sigma_{t,\tau}}\right) + \mu_{t,\tau}\sigma_{t,\tau}\phi\left(\frac{\mu_{t,\tau}}{\sigma_{t,\tau}}\right), \end{cases} \tag{13}$$

under the log-normality assumption, where $\Phi(\cdot)$ and $\phi(\cdot)$ are the standard normal cumulative distribution functions and density, respectively. An estimate of $\mu_{t,\tau}$ is obtained as the fitted value from a linear regression of returns onto the vector of predictors, while an estimate of $\sigma_{t,\tau}^2$ is obtained as the fitted value from a linear regression of the total realized variance onto the same predictors. Those estimates are further plugged into the formulas in equation (13) to obtain estimates of the physical expectations of the squared returns and truncated returns.

The specification of predictors in Z has been documented in a long list of previous literature. It is now widely accepted that models based on high frequency realized variance dominates standard

GARCH-type models (e.g., Chen and Ghysels, 2011) and thus, we follow this literature. Bekaert and Hoerova (2014) examine state-of-the-art models in the literature and consider the most general specification, where Z is a combination of a forward-looking volatility measure, the continuous variations, and the jump variations and negative returns in the last day, last week or last month:

$$\begin{aligned} RV_{t,t+\tau} &= c + \alpha VIX_t^2 + \beta^m C_{t-21,t} + \beta^w C_{t-5,t} + \beta^d C_t \\ &\quad + \gamma^m J_{t-21,t} + \gamma^w J_{t-5,t} + \gamma^d J_t \\ &\quad + \delta^m l_{t-21,t} + \delta^w l_{t-5,t} + \delta^d l_{t-1,t} + \varepsilon_{t+\tau}^{(\tau)}, \end{aligned} \tag{14}$$

where C_t and J_t are respectively continuous and discontinuous components of the daily realized variance RV_t, $C_{t-h,t}$ and $J_{t-h,t}$ respectively aggregate C_{t-j} and J_{t-j} for $j = 0, 1, \dots, h-1$, i.e. over an horizon h, and $l_{t-h,t}$ is the loss component of the return from day $t-h$ to day t. The conditional variance $\sigma_{t,\tau}^2$ is the fitted time series from the regression (14), for values of $\tau = 21, 42, \dots, 252$ days. Likewise, the conditional mean $\mu_{t,\tau}$ is the fitted time series from the regression (14) where the left-hand side is replaced by the τ-period log returns $r_{t,t+\tau}$.

Unlike the log returns and the realized variance which are closed to temporal aggregation, the quadratic payoff and its loss and gain components are not. This suggests that the term structure of physical expectations of the quadratic payoff and its components are unlikely to be a flat line unless the mean $\mu_{t,\tau}$ is negligible for all considered horizons.

2.4 Data

2.4.1 Option Data

For the estimation of the S&P 500 risk-neutral quadratic payoff, we rely on S&P 500 option prices obtained from the IvyDB OptionMetrics database for the January 1996 to December 2015 period. We exclude options with missing bid-ask prices, missing implied volatility, zero bids, negative bid-ask spreads, and options with zero open interest (e.g, Carr and Wu, 2009). Following Bakshi et al. (2003), we restrict the sample to out-of-the-money options. We further remove options with moneyness lower than 0.2 or higher than 1.8. To ensure that our results are not driven by misleading prices, we follow Conrad et al. (2013) and exclude options that do not satisfy the usual option price bounds e.g. call options with a price higher than the underlying, and options with less than 7 days to maturity.

2.4.2 Return Data

To construct the physical realized variance and perform volatility forecasts, we obtain intradaily S&P 500 cash index data spanning the period from January 1990 to December 2015 from Tick-Data.com, for a total of $6,542$ trading days. On a given day, we use the last record in each five-minute interval to build a grid of five-minute equity index log-returns. Following Andersen et al. (2001, 2003) and Barndorff-Nielsen et al. (2010), we construct the realized variance on any given trading day t, where $r_{j,t}$ is the jth five-minute log-return, and n_t is the number of (five-minute) intradaily returns recorded on that day.[3] We add the squared overnight log-return to the realized variance. The realized variance between day t and $t+\tau$ are computed by accumulating the daily realized variances.

2.5 Preliminary analysis

2.5.1 The Term Structure of the Risk-Neutral Expected Quadratic Payoff

To estimate $\mathbb{E}_t^{\mathbb{Q}}\left[l_{t,t+\tau}^2\right]$ or $\mathbb{E}_t^{\mathbb{Q}}\left[g_{t,t+\tau}^2\right]$ for each maturity τ, we use options with maturity close to τ and do interpolations.[4] In Panel A of Figure 1, we plot the time series average of risk-neutral expected quadratic payoff and its loss and gain components for maturities of 1, 3, 6, 9 and 12 months. We find that the average term structures of the risk-neutral expected quadratic payoff and its loss component are, in general, upward sloping. On the other hand, we find that the term structure of risk-neutral expected gain quadratic payoff is flat. This last result is consistent with the findings of Dew-Becker et al. (2017) who find that the term structure of the upside component of the VIX is flat. [5]

To investigate time variations in these term structures, in Panel A and B of Figure 2 we plot the 6-month (the level) and the 12- minus 2-month (the slope) for the risk-neutral expected quadratic

[3] On a typical trading day, we observe $n_t = 78$ five-minute returns.

[4] In the data, we do not always observe options with the exact maturity τ. In order to find $\mathbb{E}_t^{\mathbb{Q}}\left[l_{t,t+\tau}^2\right]$ or $\mathbb{E}_t^{\mathbb{Q}}\left[g_{t,t+\tau}^2\right]$ at the exact maturity τ we either interpolate or extrapolate to find the exact value. For example, if we wish to find $\mathbb{E}_t^{\mathbb{Q}}\left[l_{t,t+30}^2\right]$, i.e. with maturity $\tau = 30$ days, we interpolate between $\mathbb{E}_t^{\mathbb{Q}}\left[l_{t,t+\tau_1}^2\right]$ and $\mathbb{E}_t^{\mathbb{Q}}\left[l_{t,t+\tau_2}^2\right]$ to obtain $\mathbb{E}_t^{\mathbb{Q}}\left[l_{t,t+30}^2\right]$, where τ_1 is the closest observed maturity below 30 days, and τ_2 is the closest observed maturity over 30 days. In cases where we do not observe τ_2 in the data, we extrapolate τ_1 to obtain the exact maturity.

[5] Dew-Becker et al. (2017) also compute the term structure of forward variance prices as $F_t^{rv,\tau} \equiv \mathbb{E}_t^{\mathbb{Q}}\left[RV_{t+\tau-1,t+\tau}\right]$. The forward variance price $F_t^{rv,\tau}$ is essentially the month t risk-neutral expectation of realized variance from end of month $t+\tau-1$ to end of month $t+\tau$. In the Internet Appendix, we compute forward prices for the risk-neutral expected quadratic payoff and its components. In Figure A1 of the Internet Appendix, we plot the term structure of average forward prices for the quadratic payoff and its loss and gain components. We find that both the level and slope of the total quadratic payoff is similar to Dew-Becker et al. (2017). In contrast to their approach, we are also able to compute the term structure of average forward prices for the quadratic loss and gain, and find that they are in general upward sloping.

payoff and its components, respectively. We find that both the level and the slope display important time variations, and have spikes and troughs during crises. We also notice that, although the slopes are mostly positive, they are negative during crises. These observed patterns are in line with the fact that during crises investors expect a recovery in the long-run rather than the short-run.

2.5.2 The Term Structure of the Physical Expected Quadratic Payoff

In Panel B of Figure 1, we plot the time series average of the term structure of physical expected quadratic payoff and its loss and gain components for maturities of 1, 3, 6, 9 and 12 months. We find that the term structure of the expected quadratic loss is downward sloping. Since the quadratic loss is a measure of loss uncertainty, this suggests that investors face more uncertainty about losses in the short-run vs. the long-run. On the other hand, we find that the term structure of the expected quadratic gain is upward sloping. Since the expected quadratic gain is a measure of the gain uncertainty, this suggests that investors face more uncertainty about gains in the long-run vs. short-run. Comparing their term structures, we observe that the level of the expected quadratic gain dominates the level of the expected quadratic loss across all horizons, and even more so in the long-run, leading to the upward sloping pattern in the total expected quadratic payoff. The relatively larger values of the expected quadratic gain are consistent with the fact that the S&P 500 cash index has historically yielded a positive annual return of 7%.

To evaluate the time variation in these term structures, in Panel C and D of Figure 2, we plot the 6-month maturity (the level) and the 12- minus 2-month maturity (the slope) for the expected quadratic payoff and its components, respectively. As for its risk-neutral counterpart, we find substantial variations in both the level and the slope. We find that the expected quadratic payoff and expected quadratic gain have in general positive and occasionally negative slopes. On the other hand, the expected quadratic loss slopes are almost always negative and very negative during the 2008 financial crisis. These observed patterns are in line with the fact that investors expect a growth opportunity in the long-run rather than the short-run.

Further, in Figure 1, we observe a common upward sloping pattern for the term structures of the risk-neutral and physical expected quadratic payoff and its components. The only exception is the physical expected quadratic loss which is downward sloping. Bakshi et al. (2003) show that under certain conditions, the risk-neutral distribution can be obtained by exponentially tilting the real-world density, with the tilt determined by the risk-aversion of investors. This means that the observed upward sloping risk-neutral expected quadratic loss vs. the downward sloping term structure of the physical quadratic loss may be explained by investors' increasing risk-aversion as the investment horizon increases.

In Table 1, we present time series means of the risk-neutral and physical expected quadratic payoff together with their loss and gain components. For each of the mean we also report, in parentheses, Newey and West (1987) adjusted standard errors. The mean values for the risk-neutral expected quadratic payoff increase as the maturity horizon increases from 45.30 at 1 month to 49.94 at 12 months in monthly percentage squared units. The mean values for the physical expected quadratic payoff are much lower but also increase as the maturity horizon increases from 26.28 at 1 month to 28.64 at 12 months. The risk-neutral expected quadratic loss are much higher than the risk-neutral expected quadratic gain for any given horizon and the wedge is the same for different maturity horizons. For example, at 2 months, the risk-neutral expected quadratic loss is 31.16 and the risk-neutral expected quadratic gain is 14.39; the wedge is about 17 which is similar to the the wedge at 4, 6, 8 and 12 months. However, the physical expected quadratic loss is much lower than the physical expected quadratic gain for any given horizon and this wedge is increasing as the horizon increases. For example, at 2 months, the physical expected quadratic loss is 10.05 and the physical expected quadratic gain is 16.45; the wedge is about 6 and this wedge is strictly increasing to roughly 16 at 12 months. We also see that the standard errors of the means for all these quantities are decreasing in the maturity. Finally, we find that all the means are statistically different from zero.

2.5.3 The Term Structure of the Quadratic Risk Premium

Next, we turn to study the term structure of the quadratic risk premium. In Table 1, we also present time series means of the quadratic risk premium and its loss and gain components. On average, the quadratic risk premium is positive, equal to 19.04, 20.01 and 21.32 at 1, 6, 12 months, respectively. Both the loss quadratic risk premium and the gain quadratic risk premium are positive. However, the loss quadratic risk premium is dominating the gain quadratic risk premium at all horizons. For example, QRP^l is 21.98 while QRP^g is 2.91 at 3 months; QRP^l is 26.89 while QRP^g is 5.78 at 9 months. The average QRP^g is small at 1-month horizon and not statistically different from zero. In general, we observe that the standard error of the average quadratic risk premium (and its components), which represent the insurance cost (either against downside risk, upside risk or the net cost of hedging downside risk), increases with the horizon. Nevertheless, apart from the 1-month average QRP^g, we find that all means are significantly different from zero.

2.5.4 Principal Component Analysis

In general, structural and reduced form asset pricing models have a very tight factor structure, implying that different expectations (whether risk-neutral or physical) are all driven by a very low number of factors (e.g., in reduced-form option pricing models the largest number of factors considered in the literature so far is three). Nevertheless, our analysis deals with the joint term structures of two uncertainty components (loss and gain) under two different probability measures ($\mathbb{Q}$ and $\mathbb{P}$). To pin down the number of factors observed in the data, we run a principal component analysis of the term structure of four quantities: the loss and gain components of the physical and risk-neutral expected quadratic payoff. Alternatively, one can choose to use the loss and gain components of the physical expected quadratic payoff and the QRP or the loss and gain components of the risk-neutral expected quadratic payoff and the QRP. There is no difference between these three choices.

Table 2 shows the explanatory powers of the first 3 principle components. We find that the first 3 principal components are enough to explain 91.39% of the variation of the term structure of the loss and gain physical expected quadratic payoff and the QRP (there are 48 variables because we include four quantities with 12 maturities). The first principal component explains 56.76%, the second explains 26.29% and the third explains 7.98% of the variations. The immediate implication of these findings is that any model (whether reduced-form or structural) that aims to jointly fit these various terms structures should include at least three factors.

3 A Model for the Joint Term Structure of Quadratic Loss and Gain

In search of a flexible reduced-form model to accommodate different kinds of distribution asymmetry and the term structure of $\mu_2^{\mathbb{Q}-}(t,\tau)$ and $\mu_2^{\mathbb{Q}+}(t,\tau)$, we study the recent model proposed by Andersen et al. (2015). This model is ideal for our analysis (1) it is built to disentangle the dynamic of the positive and negative jumps; (2) it is a three-factor framework which would maximize the model chances of fitting the term structure of expected quadratic loss and gain and their risk premium since we find three principal components are needed to fit the targeted term structures in section 2.5.4; (3) Since it is an affine model, it is tractable and enables us to compute all the quantities of interest in closed-form. In this section, we discuss the Andersen et al. (2015) model and some variants of this three-factor model. We use the two-factor diffusion model of Christoffersen et al. (2009) as the baseline model. Finally, we introduce a set of different specifications for the pricing kernel, including the baseline specification in which jumps are not priced.

3.1 The Andersen et al. (2015)'s Risk-neutral Specification

In the three-factor jump-diffusive stochastic volatility model of Andersen et al. (2015), the underlying asset price evolves according to the following general dynamics (under $\mathbb{Q}$):

$$
\begin{aligned}
\frac{dS_t}{S_{t-}} &= (r_{f,t} - \delta_t)\,dt + \sqrt{V_{1t}}dW_{1t}^{\mathbb{Q}} + \sqrt{V_{2t}}dW_{2t}^{\mathbb{Q}} + \eta\sqrt{V_{3t}}dW_{3t}^{\mathbb{Q}} + \int_{R^2} (e^x - 1)\,\mu^{\mathbb{Q}}(dt, dx, dy) \\
dV_{1t} &= \kappa_1 (\bar{v}_1 - V_{1t})\,dt + \sigma_1\sqrt{V_{1t}}dB_{1t}^{\mathbb{Q}} + \mu_1 \int_{R^2} x^2 1_{\{x<0\}}\mu(dt, dx, dy) \\
dV_{2t} &= \kappa_2 (\bar{v}_2 - V_{2t})\,dt + \sigma_2\sqrt{V_{2t}}dB_{2t}^{\mathbb{Q}} \\
dV_{3t} &= -\kappa_3 V_{3t}dt + \mu_3 \int_{R^2} \left[(1-\rho_3)\,x^2 1_{\{x<0\}} + \rho_3 y^2\right]\mu(dt, dx, dy),
\end{aligned}
$$

where $r_{f,t}$ and δ_t refer to the instantaneous risk-free rate and the dividend yield, respectively, $\left(W_{1t}^{\mathbb{Q}}, W_{2t}^{\mathbb{Q}}, W_{3t}^{\mathbb{Q}}, B_{1t}^{\mathbb{Q}}, B_{2t}^{\mathbb{Q}}\right)$ is a five-dimensional Brownian motion with $corr\left(W_{1t}^{\mathbb{Q}}, B_{1t}^{\mathbb{Q}}\right) = \rho_1$ and $corr\left(W_{2t}^{\mathbb{Q}}, B_{2t}^{\mathbb{Q}}\right) = \rho_2$, while the remaining Brownian motions are mutually independent, $\mu^{\mathbb{Q}}(dt, dx, dy) \equiv \mu(dt, dx, dy) - \nu_t^{\mathbb{Q}}(dx, dy)\,dt$, where $\nu_t^{\mathbb{Q}}(dx, dy)$ is the risk-neutral compensator for the jump measure μ, and is assumed to be

$$
\begin{aligned}
\nu_t^{\mathbb{Q}}(dx, dy) = \Big\{ &\left(c_t^- 1_{\{x<0\}}\lambda_- e^{-\lambda_-|x|} + c_t^+ 1_{\{x>0\}}\lambda_+ e^{-\lambda_+|x|}\right) 1_{\{y=0\}} \\
&+ c_t^- 1_{\{x=0,y<0\}}\lambda_- e^{-\lambda_-|y|} \Big\}\, dx \otimes dy,
\end{aligned}
\tag{15}
$$

where time-varying negative and positive jumps are governed by distinct coefficients: c_t^- and c_t^+, respectively. These coefficients evolve as affine functions of the state vectors

$$
c_t^- = c_0^- + c_1^- V_{1t-} + c_2^- V_{2t-} + c_3^- V_{3t-}, \qquad c_t^+ = c_0^+ + c_1^+ V_{1t-} + c_2^+ V_{2t-} + c_3^+ V_{3t-}.
$$

These three factors have distinctive features: V_{2t} is a pure diffusion process, V_{3t} is a pure jump process, while innovation in V_{1t} combine a diffusion and a jump component. Furthermore, one of the key features of the AFT model is its ability to break the tight link between expected negative and positive jump variation imposed by other traditional jump diffusion models. More precisely, the AFT model implies that

$$E_t^{\mathbb{Q}}\left[NJ_{t,t+\tau}\right] = \frac{2}{\lambda_-^2\tau}E_t^{\mathbb{Q}}\left[\int_t^{t+\tau} c_s^- ds\right], \quad E_t^{\mathbb{Q}}\left[PJ_{t,t+\tau}\right] = \frac{2}{\lambda_+^2\tau}E_t^{\mathbb{Q}}\left[\int_t^{t+\tau} c_s^+ ds\right], \tag{16}$$

where $PJ_{t,t+\tau}$ and $NJ_{t,t+\tau}$ are the positive and negative jump variations between t and $t+\tau$, respectively.

To better understand the key features of this general model to match the observed term structures of risk-neutral expected quadratic loss and gain, we focus on two dimensions: (1) The number of factors. Compared to the three-factor framework, we ask if two factors are enough and which two-factor alternatives generate the best fit; (2) The model's ability to differentiate between the negative and positive jump distribution. We ask if the symmetric jump distribution can still fit the term structures. We label the unrestricted general model AFT4 and consider the following nested specifications:

- AFT0, there is no jumps. This corresponds to the two-factor diffusion model studied extensively in Christoffersen et al. (2009). This is equivalent to suppressing all the jumps related components ($\eta = 0$ and $\mu_1 = 0$) and the third factor V_{3t}.
- AFT1, there is no pure-jump process. This corresponds to suppressing V_{3t}.
- AFT2, there is no pure-diffusion process. This corresponds to suppressing V_{2t}. In this model both variance factors V_{1t} and V_{3t} jump, implying that it can be used to judge the benefit of having jumps in volatility, something that the option pricing literature has been debating quite a bit about.
- AFT3, the expected negative jump variation equals the expected positive jump variation. This corresponds to a three-factor model which assumes the same distribution for positive and negative jumps. It is equivalent to imposing that $\lambda_- = \lambda_+$, and $c_t^- = c_t^+$. The AFT3 is representative of most of existing option pricing and variance swap models as it does not differentiate between positive and negative jumps intensity (see e.g. Bates, 2012, Christoffersen et al., 2012, Eraker, 2004, Chernov et al., 2003, Huang and Wu, 2004, Amengual and Xiu, 2018, and Ait-Sahalia et al., 2015).

One interesting model variation is the 3-factor model in which $\eta = 0$. This makes V_{3t} a pure jump process in the sense that it only drives the jump intensity while not entering in the diffusive volatility.[6] We compare two 3-factor models in which $\eta = 0$ and $\eta \neq 0$ in the Internet Appendix and find that $\eta \neq 0$ is important for accurate pricing of truncated second moments.

3.2 The Radon-Nikodym Derivative

In this paper, our goal is to understand statistical properties of stock returns distribution that are essential to reproduce the observed term structures of $\mu_2^{\mathbb{Q}+}(t,\tau)$, $\mu_2^{\mathbb{Q}-}(t,\tau)$ and the stochastic discount factor specifications able to replicate the observed spreads $\mu_2^{\mathbb{Q}-}(t,\tau) - \mu_2^{\mathbb{P}-}(t,\tau)$ and $\mu_2^{\mathbb{Q}+}(t,\tau) - \mu_2^{\mathbb{P}+}(t,\tau)$. To do that, we need to specify a Radon-Nikodym Derivative (the law of change of measure). We specify the most flexible Radon-Nikodym derivative preserving the same model structure under the physical dynamic. Our Radon-Nikodym derivative is the product of the two derivatives separately governing the compensation of continuous variations and jump variations:

$$\left(\frac{dQ}{dP}\right)_t = \left(\frac{dQ}{dP}\right)_t^c \left(\frac{dQ}{dP}\right)_t^j,$$

where

$$\left(\frac{dQ}{dP}\right)_t^c = \exp\left\{\int_0^t \theta_s^{r\top} dW_s^{\mathbb{P}} + \int_0^t \bar{\theta}_s^{v\top} d\bar{B}_s^{\mathbb{P}} - \frac{1}{2}\int_0^t \left(\theta_s^{r\top}\theta_s^r + \bar{\theta}_s^{v\top}\bar{\theta}_s^v\right) ds\right\},$$

and

$$\left(\frac{dQ}{dP}\right)_t^j = \mathcal{E}\left(\int_0^t \int_{R^2} \Psi_s(x,y)\,\mu^{\mathbb{P}}(ds,dx,dy)\right),$$

with $\mathcal{E}$ referring to the stochastic exponential, $dW_{jt}^{\mathbb{P}} \equiv dW_{jt}^{\mathbb{Q}} + \theta_t^r(j)\,dt$, $d\bar{B}_{jt}^{\mathbb{P}} \equiv d\bar{B}_{jt}^{\mathbb{Q}} + \bar{\theta}_t^v(j)\,dt$, $dB_{jt}^{\mathbb{Q}} = \rho_j dW_{jt}^{\mathbb{Q}} + \sqrt{1-\rho_j^2}\,d\bar{B}_{jt}^{\mathbb{Q}}$ and $\mu^{\mathbb{P}}(dt,dx,dy) = \mu(dt,dx,dy) - \nu_t^{\mathbb{P}}(dx,dy)\,dt$.

With the appropriate choice of the price of risk parameters θ_t^r, $\bar{\theta}_t^v$ and the physical compensator $\nu_t^{\mathbb{P}}(dx,dy)$, we can show that the resulting physical dynamic preserves the exact structure as the risk-neutral dynamic. In particular, the price of jump risk, $\Psi_t(x,y)$, is given by:

$$\Psi_t(x,y) \equiv \frac{\nu_t^{\mathbb{Q}}(dx,dy)}{\nu_t^{\mathbb{P}}(dx,dy)} - 1,$$

where

$$\frac{\nu_t^{\mathbb{Q}}(dx,dy)}{\nu_t^{\mathbb{P}}(dx,dy)} = \begin{cases} \frac{c_t^+}{c_t^{\mathbb{P}+}}\frac{\lambda_+}{\lambda_+^{\mathbb{P}}}\exp\left(-\left(\lambda_+ - \lambda_+^{\mathbb{P}}\right)x\right) & x>0 \quad y=0 \\ \frac{c_t^-}{c_t^{\mathbb{P}-}}\frac{\lambda_-}{\lambda_-^{\mathbb{P}}}\exp\left(\left(\lambda_- - \lambda_-^{\mathbb{P}}\right)x\right) & x<0 \quad y=0 \\ \frac{c_t^-}{c_t^{\mathbb{P}-}}\frac{\lambda_-}{\lambda_-^{\mathbb{P}}}\exp\left(\left(\lambda_- - \lambda_-^{\mathbb{P}}\right)y\right) & x=0 \quad y<0 \end{cases}.$$

[6]Several contributions including Santa-Clara and Yan (2010), Christoffersen et al. (2012), and Andersen et al. (2015) find evidence for a pure jump component in the pricing of S&P500 options

Is the premium inherent in hedging bad shocks substantially different from the one required to be exposed to good shocks? The evidence presented in Section 2 overwhelmingly points to two very different premia. Another more challenging question is if we need to specify a maximum flexible pricing kernel where all the parameters for jumps intensities are shifted from $\mathbb{Q}$-measure to $\mathbb{P}$-measure by different amounts, or if there is a parsimonious specification (imposes more restrictions between the $\mathbb{Q}$- and $\mathbb{P}$-dynamics) that is able to simultaneously replicate the observed dynamic of the term structures of the loss and gain QRP. To shed light on these issues, in the estimation investigation we distinguish between the following restrictions on the Radon-Nikodym derivative (where the unrestricted specification is labeled RND4):

1. RND0: Jumps are not priced, this is equivalent to imposing $c_j^{\mathbb{P}+} = c_j^+$, $c_j^{\mathbb{P}-} = c_j^-$, for $j = 0, 1, 2, 3$ $\lambda_-^{\mathbb{P}} = \lambda_-$ and $\lambda_+^{\mathbb{P}} = \lambda_+$. Note that this is the equivalent of setting $\Psi_t(x, y) = 0$, or equivalently $\left(\frac{dQ}{dP}\right)_t^j = 1$.
2. RND1: The price of positive jumps equals the price of negative jumps, or more formally $\Psi_t(x, y)$ is independent of the sign of x. Note that this is the implicit restriction imposed by traditional affine jump diffusion option pricing models, e.g. Eraker (2004),Santa-Clara and Yan (2010), Christoffersen et al. (2012) and Bates (2012).
3. RND2: Negative jumps are not priced $\Longleftrightarrow \lambda_-^{\mathbb{P}} = \lambda_-$, $c_j^{\mathbb{P}-} = c_j^-$, for $j = 0, 1, 2, 3,$.
4. RND3: Positive jumps are not priced $\Longleftrightarrow \lambda_+^{\mathbb{P}} = \lambda_+$, $c_j^{\mathbb{P}+} = c_j^+$, for $j = 0, 1, 2, 3,$.

4 Estimation

We largely rely on the recent paper Feunou and Okou (2018) which proposes to estimate affine option pricing models using risk-neutral moments instead of raw option prices. Unlike option prices, cumulants (central moments) are linear functions of unobserved factors. Hence using cumulants enables us to circumvent major challenges usually encountered in the estimation of latent factor option pricing models.

Given that the AFT model is affine, the linear Kalman filter appears as a natural estimation technique. The AFT model can easily be casted in a (linear) state-space form where the measurement equations relate the observed or model-free risk-neutral cumulants to the latent factors (state variables), and the transition equations describe the dynamic of these factors. However, unlike the setup in Feunou and Okou (2018), we are mainly interested in the term structures of expected quadratic loss and gain which turn out to be non-linear functions of the factors. Hence we will have two sets of measurement equations: (1) linear, which relate the risk-neutral variances and third order cumulants to the factors; (2) non-linear, which relate the risk-neutral expected quadratic loss and gain to the factors. We will use only the first set of measurement equations in the linear Kalman filtering step, and conditional on the filtered factors we will compute the likelihood of the risk-neutral expected quadratic loss and gain.

4.1 Risk-neutral Cumulants Likelihood

On a given day t, we stack together the n^{th}-order risk-neutral cumulant observed at distinct maturities in a vector denoted by $CUM_t^{(n)\mathbb{Q}} = (CUM_{t,\tau_1}^{(n)\mathbb{Q}}, ..., CUM_{t,\tau_J}^{(n)\mathbb{Q}})^\top$, where $n \in \{2, 3\}$. We further stack the second and third cumulant vector in $CUM_t^{\mathbb{Q}} = \left(CUM_t^{(2)\mathbb{Q}\top}, CUM_t^{(3)\mathbb{Q}\top}\right)^\top$ to build a $2J \times 1$ vector. This implies the following linear measurement equation:

$$CUM_t^{\mathbb{Q}} = \Gamma_0^{cum} + \Gamma_1^{cum} V_t + \Omega_{cum}^{1/2} \vartheta_t^{cum}, \tag{17}$$

where the dimension of the unobserved state vector (V_t) is 3. Notably, Γ_0^{cum} and Γ_1^{cum} are $2J \times 1$ and $2J \times 3$ matrices of coefficients whose analytical expressions depend explicitly on Q-parameters as shown in Feunou and Okou (2018). The last term in Equation 17 is a vector of observation errors, where Ω_{cum} is a $2J \times 2J$ diagonal covariance matrix, and ϑ_t^{cum} denotes a $2J \times 1$ vector of independent and identically distributed (i.i.d.) standard Gaussian disturbances.

As shown in Feunou and Okou (2018), the transition equations for the three factors in the AFT model is:

$$V_{t+1} = \Phi_0 + \Phi_1 V_t + \varepsilon_{t+1}, \tag{18}$$

where

$$\begin{aligned}
\Phi_0 &\equiv \Delta t K_0^{\mathbb{P}}, \quad K_0^{\mathbb{P}} = \begin{pmatrix} \kappa_1^{\mathbb{P}} \bar{v}_1^{\mathbb{P}} + \mu_1 \bar{\lambda}_-^{\mathbb{P}} c_0^{\mathbb{P}-} \\ \kappa_2^{\mathbb{P}} \bar{v}_2^{\mathbb{P}} \\ \mu_3 \bar{\lambda}_-^{\mathbb{P}} c_0^{\mathbb{P}-} \end{pmatrix} \\
\Phi_1 &\equiv I_3 + \Delta_t K_1^{\mathbb{P}}, \quad K_1^{\mathbb{P}} = \begin{bmatrix} -\kappa_1^{\mathbb{P}} + \mu_1 \bar{\lambda}_-^{\mathbb{P}} c_1^{\mathbb{P}-} & \mu_1 \bar{\lambda}_-^{\mathbb{P}} c_2^{\mathbb{P}-} & \mu_1 \bar{\lambda}_-^{\mathbb{P}} c_3^{\mathbb{P}-} \\ 0 & -\kappa_2^{\mathbb{P}} & 0 \\ \mu_3 \bar{\lambda}_-^{\mathbb{P}} c_1^{\mathbb{P}-} & \mu_3 \bar{\lambda}_-^{\mathbb{P}} c_2^{\mathbb{P}-} & -\kappa_3^{\mathbb{P}} + \mu_3 \bar{\lambda}_-^{\mathbb{P}} c_3^{\mathbb{P}-} \end{bmatrix},
\end{aligned}$$

I_3 is a 3×3 identity matrix, $\bar{\lambda}_{-}^{\mathbb{P}} = 2/\left(\lambda_{-}^{\mathbb{P}}\right)^2$, and Δt is set to $1/252$ to reflect a daily time step. Moreover, the transition noise is $\varepsilon_{t+1} \equiv \left(\varepsilon_{1t+1}, \varepsilon_{2t+1}, \varepsilon_{3t+1}\right)^\top$, with a conditional covariance matrix $Var_t^{\mathbb{P}}\left(\varepsilon_{t+1}\right) = \Delta t \Sigma\left(V_t\right)$, where

$$\Sigma\left(V_t\right) = \begin{bmatrix} \sigma_1^2 V_{1t} + \mu_1^2 \lambda_{-}^{\mathbb{P}*} c_t^{\mathbb{P}-} & 0 & \mu_1 \mu_3 \left(1-\rho_3\right) \lambda_{-}^{\mathbb{P}*} c_t^{\mathbb{P}-} \\ 0 & \sigma_2^2 V_{2t} & 0 \\ \mu_1 \mu_3 \left(1-\rho_3\right) \lambda_{-}^{\mathbb{P}*} c_t^{\mathbb{P}-} & 0 & \mu_3^2 \left[\left(1-\rho_3\right)^2 + \rho_3^2\right] \lambda_{-}^{\mathbb{P}*} c_t^{\mathbb{P}-} \end{bmatrix},$$

where $\lambda_{-}^{\mathbb{P}*} = 24/\left(\lambda_{-}^{\mathbb{P}}\right)^4$.

The system (17)-(18) gives the state-space representation of the AFT model. The marginal moments (mean and variance) of the latent vector are used to initialize the filter, by setting $V_{0|0} = -\left(K_1^{\mathbb{P}}\right)^{-1} K_0^{\mathbb{P}}$, and $vec\left(P_{0|0}\right) = \Delta t \left(I_9 - \Phi_1 \otimes \Phi_1\right)^{-1} vec\left(\Sigma\left(V_{0|0}\right)\right)$, where I_9 is a 9×9 identity matrix, and $\otimes$ is the Kronecker product. Now, consider that $V_{t|t}$ and $P_{t|t}$ are available at a generic iteration t. Then, the filter proceeds recursively through the forecasting step:

$$\begin{cases} V_{t+1|t} = \Phi_0 + \Phi_1 V_{t|t} \\ P_{t+1|t} = \Phi_1 P_{t|t} \Phi_1^\top + \Delta t \Sigma\left(V_{t|t}\right) \\ CUM_{t+1|t}^{\mathbb{Q}} = \Gamma_0 + \Gamma_1 V_{t+1|t} \\ M_{t+1|t} = \Gamma_1 P_{t+1|t} \Gamma_1^\top + \Omega_{cum}, \end{cases} \tag{19}$$

and the updating step:

$$\begin{cases} V_{t+1|t+1} = \left[V_{t+1|t} + P_{t+1|t} \Gamma_1^\top M_{t+1|t}^{-1} \left(CUM_{t+1}^{\mathbb{Q}} - CUM_{t+1|t}^{\mathbb{Q}}\right)\right]_+, \\ P_{t+1|t+1} = P_{t+1|t} - P_{t+1|t} \Gamma_1^\top M_{t+1|t}^{-1} \Gamma_1 P_{t+1|t}, \end{cases} \tag{20}$$

where $[V]_+$ returns a vector whose i^{th} element is $\max\left(V_i, 0\right)$. This additional condition ensures that latent factor estimates remain positive for all iterations — a crucial property for stochastic volatility factors that cannot assume negative values. Finally, we construct a Gaussian quasi log-likelihood for the cumulants:

$$Lik^{CUM} = -\frac{1}{2} \sum_{t=1}^{T} \left[\ln\left(\left(2\pi\right)^{2J} \det\left(M_{t|t-1}\right)\right) + \xi_{t,cum}^\top M_{t|t-1}^{-1} \xi_{t,cum}\right], \tag{21}$$

where $\xi_{t,cum} \equiv CUM_t^{\mathbb{Q}} - CUM_{t|t-1}^{\mathbb{Q}}$.

4.2 Risk-neutral Expected Quadratic Loss and Gain Likelihood

We use equation (5) to compute model implied $\mu_2^{\mathbb{Q}+}(t,\tau)$ and $\mu_2^{\mathbb{Q}-}(t,\tau)$. Note that $\mathbb{E}_t^{\mathbb{Q}}\left[r_{t,t+\tau}^2\right] = CUM_{t,\tau}^{(2)\mathbb{Q}} + \left(CUM_{t,\tau}^{(1)\mathbb{Q}}\right)^2$ and both $CUM_{t,\tau}^{(2)\mathbb{Q}}$ and $CUM_{t,\tau}^{(1)\mathbb{Q}}$ are computed analytically within the AFT framework following Feunou and Okou (2018). We follow Fang and Oosterlee (2008) and compute $\Lambda^{\mathbb{Q}}(t,\tau)$ as:

$$\begin{aligned}
\Lambda^{\mathbb{Q}}(t,\tau) &= \frac{2}{\pi}\int_0^{+\infty}\frac{Im\left(\varphi_{t,\tau}^{(2)}(-iv)\right)}{v}dv \equiv \mu_2^{\mathbb{Q}-}(t,\tau) - \mu_2^{\mathbb{Q}+}(t,\tau) \\
&= \mathbb{E}_t^{\mathbb{Q}}\left[r_{t,t+\tau}^2\left\{\mathbb{I}\left(r_{t,t+\tau}<0\right) - \mathbb{I}\left(r_{t,t+\tau}>0\right)\right\}\right] = \sum_{k=0}^{N-1} Re\left\{\varphi_{t,\tau}\left(i\frac{k\pi}{b-a}\right)\exp\left(-i\frac{k\pi a}{b-a}\right)\right\}\omega_k,
\end{aligned}$$

where i stands for the imaginary unit, $\mathbb{I}$ the indicator function,

$$\omega_0 \equiv \frac{1}{b-a}\int_a^b y^2\left\{\mathbb{I}(y<0) - \mathbb{I}(y>0)\right\}dy = \begin{cases} -\frac{b^3-a^3}{3(b-a)} & if \quad a \geq 0 \\ -\frac{b^3+a^3}{3(b-a)} & if \quad a < 0 \end{cases},$$

and for $k > 0$

$$\begin{aligned}
\omega_k &\equiv \frac{2}{b-a}\int_a^b y^2\left\{\mathbb{I}(y<0) - \mathbb{I}(y>0)\right\}\cos\left(k\pi\frac{y-a}{b-a}\right)dy \\
&= \begin{cases} -\frac{4(b-a)\left(b(-1)^k-a\right)}{(k\pi)^2} & if \quad a \geq 0 \\ -\frac{4(b-a)\left(b(-1)^k+a\right)}{(k\pi)^2} + \frac{8(b-a)^2}{(k\pi)^3}\sin\left(k\pi\frac{a}{b-a}\right) & if \quad a < 0 \end{cases}.
\end{aligned}$$

In the implementation phase, we set $a = \ln(0.01)$, $b = -a$, and $N = 100$. Recall that within the AFT framework, the risk-neutral moment generation function $\varphi_{t,\tau}(\cdot)$ is an exponential linear function of the factor V_t.[7] Hence, both $\mu_2^{\mathbb{Q}+}(t,\tau)$ and $\mu_2^{\mathbb{Q}-}(t,\tau)$ are non-linear functions of the factor V_t:

$$\begin{aligned}
\mu_2^{\mathbb{Q}+}(t,\tau) &= \mu_{2,\tau}^{\mathbb{Q}+}(V_t), \; \mu_2^{\mathbb{Q}-}(t,\tau) = \mu_{2,\tau}^{\mathbb{Q}-}(V_t) \\
Tmom_t^+ &= \left(\mu_{2,\tau_1}^{\mathbb{Q}+}, \cdots, \mu_{2,\tau_J}^{\mathbb{Q}+}\right)^\top, \; Tmom_t^- = \left(\mu_{2,\tau_1}^{\mathbb{Q}-}, \cdots, \mu_{2,\tau_J}^{\mathbb{Q}-}\right)^\top \\
Tmom_t &= \left(Tmom_t^{+\top}, Tmom_t^{-\top}\right)^\top.
\end{aligned}$$

[7] The coefficients relating $\ln(\varphi_{t,\tau}(\cdot))$ are a solution to Ordinary Differential Equations (ODEs) that can only be solved numerically (see the Internet Appendix of Andersen et al., 2015 for more details). We thank Nicola Fusari for sharing the estimation code.

We construct a Gaussian quasi log-likelihood for the truncated moments $Tmom_t$

$$Lik^{Tmom} = -\frac{1}{2}\sum_{t=1}^{T}\left[\ln\left((2\pi)^{2J}\det\left(\Omega_{Tmom}\right)\right) + \xi_{t,Tmom}^{\top}\Omega_{Tmom}^{-1}\xi_{t,Tmom}\right], \tag{22}$$

where $\xi_{t,Tmom} = Tmom_t^{(Obs)} - Tmom_t\left(V_{t|t}\right)$, Ω_{Tmom} denotes the measurement error variance, $Tmom_t^{(Obs)}$ is the time t observed risk-neutral truncated moments (computed model-free using (10)), and $V_{t|t}$ is obtained through the filtering procedure (see Equation (19) and (20)). Different models' parameters are estimated via a maximisation of $Lik^{CUM} + Lik^{Tmom}$.

4.3 Discussions

The Kalman filter is not optimal in this case, given the heteroscedasticity and non-Normality of factors and some nonlinear measurement equations. With respect to the problematic of heteroscedastic and non-Gaussian factors, we refer readers to Monfort et al. (2017) and Duan and Simonato (1999) for extensive discussions and Monte Carlo analyses suggesting that the lost of optimality is very minimal. Regarding nonlinear measurement equations, two other alternatives could have been considered: (1) locally linearize the nonlinear measurement equation: this is known as the Extended Kalman filter, or (2) use a deterministic sampling technique (known as the unscented transformation) to accurately estimate the true mean and covariance: this is known as the Unscented Kalman filter.

While in theory the Extended Kalman filter is appealing, as it allows the nonlinear measurement equation to impact directly the filtered factors, its implementation is very complex in our case. Mainly because of the nonlinear measurements themselves are approximate (we rely on approximation techniques developed in Fang and Oosterlee (2008)). In addition, this technique requires to explicitly calculate Jacobians, which for complex functions can be a difficult task in itself (i.e., requiring complicated derivatives if done analytically or being computationally costly if done numerically), if not impossible (if those functions are not differentiable).

Face with this uncertainty about the nature of the non-linearity, thus with an almost certain lack of robustness, we have traded the potential gain of optimality with a more robust and reliable approach. Let us emphasize that, nonlinear equations are used for parameters estimations, which implies that they impact the filter albeit indirectly. Finally, the computational hurdle is the major obstacle preventing us from contemplating the Unscented Kalman filter.

5 Results

In this section, we evaluate the ability of different models to fit the term structure of the expected quadratic payoff and its loss and gain components. We use Christoffersen et al. (2009) (AFT0) as our baseline model. We compare this baseline model with two other two-factor alternatives AFT1 and AFT2 and two more three-factor models AFT3 with symmetric jump distribution and AFT4 in Andersen et al. (2015). Finally, we evaluate the ability of different pricing kernels with various flexibility in fitting the QRP and its loss and gain components.

5.1 Fitting the Risk-neutral Expectations

We examine the performance of different models by relying on the root-mean-squared error:

$$RMSE \equiv \sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(Mom_t^{Mkt} - Mom_t^{Mod}\right)^2},$$

where Mom_t^{Mkt} is the time t observed risk-neutral moment and Mom_t^{Mod} is the model-implied equivalent. Results are reported in Table 3 where several conclusions can be drawn[8]. Overall, regarding the fitting of the term structure of the risk-neutral expected gain and loss, the benchmark two-factors diffusion model (AFT0) is outperformed by all the other variants.

With respect to the risk neutral quadratic loss fit, the AFT0 model's RMSE increases with horizon and ranges from 1% at two months to 2.15% at one year. The average RMSE is 1.73% which is far higher than other models' RMSEs. The best performer is the AFT4 model with an average error of 0.77% which offers approximately 56% improvement over the benchmark AFT0 model. This performance of the AFT4 model is robust across horizon, with a RMSE as low as 0.45% around horizons 4 to 5 months, which is an improvement of nearly 75% over the benchmark AFT0 model. The three-factor models (AFT3 and AFT4) outperform the other two-factor models (AFT1 and AFT2). The AFT4 model offers an improvement of approximately 15% over the AFT3 model, which underscores the importance of accounting for asymmetry in the jump distribution.

Turning to the risk-neutral quadratic gain fit, the AFT0 model's average RMSE is 2.19%, which is roughly 50% higher than other variants RMSEs. The best performing model on this front is the AFT1 model with an average RMSE of 0.98%, while the performances of the AFT2, AFT3 and AFT4 models are similar. However, the AFT0 model fits the term structure of the total risk-neutral

[8] To save space, we report risk-neutral parameter estimates in the Internet Appendix

quadratic payoff remarkably well with an average RMSE of 0.44%, which confirms the findings of Christoffersen et al. (2009). The best performer for the term structure of the total quadratic payoff is again the AFT4 model with a RMSE of about 0.19%, which offers an improvement of about 57% over the benchmarks AFT0 and AFT3. In accordance with our findings regarding the quadratic loss, this result highlights the importance of asymmetry in the jump distribution for fitting the term structure of risk-neutral variance.

On the term structure of risk-neutral skewness dimension, the benchmark AFT0 is the worst performer with an average RMSE of 0.85, whereas all the other variants have similar fit, with an average RMSE of approximately 0.15, which is almost 80% improvement over the benchmark the benchmark AFT0. This results underscores the importance of jumps when fitting the term structure of risk-neutral skewness. To better understand our findings, we plot in Figure 3 (the first two rows) the observed and models implied average term structure of risk-neutral moments. The AFT0 model is clearly unable to fit the average term structure of risk-neutral expected quadratic gain or loss. It overestimates the risk-neutral expected quadratic gain and underestimates the risk-neutral expected quadratic loss, which explains why it is able to fit the term structure of the total risk-neutral expected quadratic payoff well. Not surprisingly, the AFT0 model is outperformed by all the other variants when it comes to fitting the term structure of skewness. The most likely explanation is that jumps are essential to generate skewness, only accounting for the leverage effect is not enough.

The ranking between two-factor models (AFT1 and AFT2) is mixed. The model without pure jump process (AFT1) dominates the one without pure diffusion process (AFT2) when fitting the term structure of the risk-neutral expected quadratic loss and gain in the short end. However, this result is reversed in the long end. Figure 3 shows that, on average, the AFT1 model fits the term structure of the risk-neutral expected quadratic gain remarkably well, while the AFT2 model fits the term structure of risk-neutral expected quadratic loss very well. These results suggest that incorporating a pure jump process and having jumps in the volatility are essential for the distribution of the loss uncertainty, while a pure diffusion process is a key ingredient for the distribution of the gain uncertainty.

Both of the two-factor variants are outperformed by the most general specification (AFT4) which overall is able to reproduce the term structure of the truncated and total risk-neutral moments remarkably well. Comparing the two three-factors models (AFT3 and AFT4), we evaluate the importance of introducing a wedge between the negative and positive jump distribution. The results are mixed on this front. For the term structure of the risk-neutral expected quadratic loss and the risk-neutral expected quadratic payoff, the AFT4 model clearly outperforms the AFT3 model (which has no asymmetry in the jump distribution). However, there is no clear winner for the gain uncertainty and skewness. The AFT4 model fits better the short-end of the risk-neutral expected quadratic gain, while the AFT3 is preferred on the long-end.

5.2 Fitting the QRPs

We focus on the most flexible specification of the AFT model (AFT4) and evaluate the fitting ability of different pricing kernel specifications discussed in section 3.2. The last row of Figure 3 plots the observed and models' implied average term structure of the quadratic risk premium. It is readily apparent that the most flexible Radon-Nikodym derivative ($RND4$) is the only one which is able to fit adequately the average term structure of the quadratic risk premium and its loss and gain components. The worst performer is $RND0$ which assumes that jumps are not priced. Not pricing jumps generates a negative average term structure of the net and loss quadratic risk premium. Not pricing either positive jumps ($RND2$) or negative jumps ($RND3$) is also strongly rejected. Finally, even though a symmetric Radon-Nikodym derivative ($RND1$) which gives the same price to both the positive and negative jumps, is able to replicate the positive sign for all the three term structures, it fells short to capture the right level. Overall, it is imperative to price jumps asymmetrically in the pricing kernel.

To confirm these visual findings, we report the root mean squared error in Table 4. In addition to the Loss and Gain QRP RMSEs reported in the top panels, we have also reported the total QRP (denote Net QRP in table, because it is the difference between Loss and Gain QRP) in the bottom left panel and the skewness risk-premium (which is defined as the sum of the Loss and Gain QRP) in the bottom right panel. The numbers are roughly in line with the Figure 3's visual findings. Except for the very short maturity (3 month), $RND4$ yields the smallest RMSE across maturities and for different types of risk premium. $RND0$ is the worst performer, which implies that pricing jumps is important for the dynamic of the quadratic risk premium and its components. The average RMSE for the $RND4$ model is 3.5%, 1.2%, 1.4% and 3.2% for the Loss, gain, net and the sum QRP respectively. These numbers are substantial improvements over the benchmarks $RND0$ (which assumes that jumps are not priced) and $RND1$ (which gives the same price to both the positive and negative jumps). To be more specific, on one hand the $RND4$ model offers approximately 70%, 30%, 80% and 72% improvements over $RND0$ for the fitting of the Loss, gain, net and the sum QRP respectively. On the other hand, the $RND4$ model offers approximately 62%, 33%, 75% and 66% improvement over $RND1$ for the fitting of the Loss, gain, net and the sum QRP respectively.

We can further scrutinize the overall performance results by maturity. Table 4 reveals that the superiority of the $RND4$ pricing kernel holds across the maturity spectrum. For the Loss, net and sum QRP, the relative improvement increases with the maturity, and reaches 80% at the one year horizon.

6 Conclusion

In this paper we investigate how the amount of money paid by investors to hedge negative spikes in the stock market changes with the investment horizon. For this purpose, we estimate the quadratic payoff and its loss and gain components across time and horizon. We uncover new empirical facts which challenge most of the existing option and variance swaps pricing models. Among these facts, we find an average upward sloping term structure for the risk-neutral expected quadratic payoff and its components. We also find upward sloping term structures for the physical expected quadratic payoff and quadratic gain but a downward sloping term structure for the physical expected quadratic loss. There is significant time variation in the slopes of these term structures, and we observe that they are negative and spike during financial downturns. Finally, we find that at least three principal components are required to explain the cross-section (across maturity or horizon) of the risk-neutral and physical expected quadratic payoff and its components.

To replicate these empirical facts we focus on the Andersen et al. (2015) model and some of its restricted variants. This model is particularly appealing as it completely disentangles the dynamics of negative and positive jumps. In addition, the model has three factors which is an essential ingredient as suggested by our principal component analysis. We find that models without an asymmetric treatment of positive and negative jumps are overall rejected as they are unable to fit the term structure of the risk-neutral expected quadratic loss. Notably, this category of models covers most of the existing option and variance swap pricing models found in the literature. We also evaluate different pricing kernel specifications and find that disentangling the price of negative jumps from its positive counterpart is essential for replicating the observed term structures of loss and gain quadratic risk premium.

References

Ait-Sahalia, Yacine, Karaman, Mustafa, and Mancini, Loriano (2015), "The term structure of variance swaps and risk premia," *Working Paper, Princeton University, University of Zurich and University of Lugano* .

Amengual, Dante, and Xiu, Dacheng (2018), "Resolution of policy uncertainty and sudden declines in volatility," *Journal of Econometrics,* 203(2), 297 – 315.

Andersen, Torben, Fusari, Nicola, and Todorov, Viktor (2015), "The risk premia embedded in index options," *Journal of Financial Economics,* 117(3), 558–584.

Andersen, Torben G., Bollerslev, Tim, Diebold, Francis X., and Ebens, Heiko (2001), "The distribution of realized stock return volatility," *Journal of Financial Economics,* 61, 43–76.

Andersen, Torben G., Bollerslev, Tim, Diebold, Francis X., and Labys, Paul (2003), "Modeling and forecasting realized volatility," *Econometrica,* 71(2), 579–625.

Bakshi, G., and Madan, D. (2000), "Spanning and Derivative-security Valuation," *Journal of Financial Economics,* 55, 205–238.

Bakshi, Gurdip, Kapadia, N., and Madan, Dilip (2003), "Stock return characteristics, skew laws and the differential pricing of individual equity options," *Review of Financial Studies,* 16(1), 101–143.

Barndorff-Nielsen, Ole E., Kinnebrock, Silja, and Shephard, Neil (2010) *"Measuring downside risk: realised semivariance," in Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle,* eds. T. Bollerslev, and J. Russell, and M. Watson, Oxford: Oxford University Press, pp. 117-136.

Bates, David S. (2012), "U.s. stock market crash risk, 19262010," *Journal of Financial Economics,* 105(2), 229 – 259.

Bekaert, G., Engstrom, E., and Ermolov, A. (2015), "Bad Environments, Good Environments: A Non-Gaussian Asymmetric Volatility Model," *Journal of Econometrics,* 186(1), 258–275.

Bekaert, Geert, and Hoerova, Marie (2014), "The vix, the variance premium and stock market volatility," *Journal of Econometrics,* 183(2), 181 – 192. Analysis of Financial Data.

Bernardo, Antonio E., and Ledoit, Olivier (2000), "Gain, loss, and asset pricing," *Journal of Political Economy,* 108(1), 144–172.

Black, F., and Scholes, M. (1973), "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy,* 81(3), 637–654.

Carr, Peter, and Wu, Liuren (2009), "Variance risk premiums," *Review Financial Studies,* 22, 1311–1341.

Chen, Xilong, and Ghysels, Eric (2011), "News good or bad and its impact on volatility predictions over multiple horizons," *The Review of Financial Studies,* 24(1), 46–81.

Chernov, Mikhail, Gallant, Ronald, Ghysels, Eric, and Tauchen, George (2003), "Alternative models for stock price dynamics," *Journal of Econometrics,* 116, 225–257.

Christoffersen, Peter, Heston, Steven, and Jacobs, Kris (2009), "The shape and term structure of the index option smirk: Why multifactor stochastic volatility models work so well," *Management Science,* 55(12), 1914–1932.

Christoffersen, Peter, Jacobs, Kris, and Ornthanalai, Chayawat (2012), "Dynamic jump intensities and risk premiums: Evidence from s&p500 returns and options," *Journal of Financial Economics,* 106(3), 447 – 472.

Conrad, Jennifer, Dittmar, Robert F., and Ghysels, Eric (2013), "Ex ante skewness and expected stock returns," *The Journal of Finance,* 68(1), 85–124.

Dew-Becker, Ian, Giglio, Stefano, Le, Anh, and Rodriguez, Marius (2017), "The price of variance risk," *Journal of Financial Economics,* 123(2), 225 – 250.

Duan, Jin-Chuan, and Simonato, Jean-Guy (1999), "Estimating exponential-affine term structure models by kalman filter," *Review of Quantitative Finance and Accounting,* 13(2), 111–135.

Duffie, Darrell, Pan, Jun, and Singleton, Kenneth (2000), "Transform analysis and option pricing for affine jump-diffusions," *Econometrica,* 68, 1343–1377.

Eraker, Bjørn (2004), "Do stock prices and volatility jump? Reconciling evidence from spot and option prices" *Journal of Finance,* 59(3), 1367–1404.

Fang, Fang, and Oosterlee, Cornelis (2008), "A novel pricing method for european options based on fourier-cosine series expansions," *SIAM J. Scientific Computing,* 31, 826–848.

Feunou, Bruno, and Okou, Cédric (2018), "Risk-neutral moment-based estimation of affine option pricing models," *Journal of Applied Econometrics,* 33(7), 1007–1025.

Feunou, Bruno, Lopez Aliouchkin, Ricardo, Tédongap, Roméo, and Xu, Lai (2019), "Loss uncertainty, gain uncertainty, and expected stock returns," *Working Paper, Bank of Canada, Syracuse University and ESSEC Business School.*

Huang, Jing-Zhi, and Wu, Liuren (2004), "Specification analysis of option pricing models based on time-changed lévy processes," *Journal of Finance,* 59(3), 1405–1439.

Jiang, G. J., and Tian, Y. S. (2005), "The Model-Free Implied Volatility and Its Information Content," *Review of Financial Studies,* 18(4), 1305–1342.

Kilic, Mete, and Shaliastovich, Ivan (2019), "Good and bad variance premia and expected returns," *Management Science,* 65(6), 2522–2544.

Monfort, Alain, Pegoraro, Fulvio, Renne, Jean-Paul, and Roussellet, Guillaume (2017), "Staying at zero with affine processes: An application to term structure modelling," *Journal of Econometrics,* 201(2), 348 – 366. THEORETICAL AND FINANCIAL ECONOMETRICS: ESSAYS IN HONOR OF C. GOURIEROUX.

Newey, W.K., and West, K.D. (1987), "A Simple, Positive Semi-definite Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica,* 55, 703–708.

Patton, A., and Sheppard, K. (2015), "Good Volatility, Bad Volatility: Signed Jumps and the Persistence of Volatility," *Review of Economics and Statistics,* 97(3), 683–697.

Santa-Clara, Pedro, and Yan, Shu (2010), "Crashes, volatility, and the equity premium: Lessons from s&p 500 options," *The Review of Economics and Statistics,* 92(2), 435–451.

Table 1: **Descriptive Statistics**

In this table, we report the time-series mean for all maturities ranging from 1 to 12 months for a set of variables that includes: the risk-neutral expected quadratic payoff ($\mathbb{E}^{\mathbb{Q}}\left[r^2\right]$, $\mathbb{E}^{\mathbb{Q}}\left[l^2\right]$, $\mathbb{E}^{\mathbb{Q}}\left[g^2\right]$), the expected quadratic payoff ($\mathbb{E}\left[r^2\right]$, $\mathbb{E}\left[l^2\right]$, $\mathbb{E}\left[g^2\right]$), and the quadratic risk premium (QRP, QRP^l, QRP^g). Below each mean, in parentheses we also report the Newey and West (1987) standard error. All statistics are monthly squared percentage values. The sample period is from January 1996 to December 2015.

	Mean											
Maturity	1	2	3	4	5	6	7	8	9	10	11	12
$\mathbb{E}\left[r^2\right]$	26.28	26.49	27.42	28.20	27.97	27.88	27.97	28.06	28.06	28.16	28.38	28.64
	(2.49)	(2.07)	(2.07)	(2.01)	(1.66)	(1.45)	(1.34)	(1.25)	(1.16)	(1.09)	(1.04)	(1.00)
$\mathbb{E}\left[l^2\right]$	10.87	10.05	10.05	9.95	9.09	8.35	7.90	7.43	6.95	6.56	6.32	6.14
	(1.58)	(1.30)	(1.54)	(1.54)	(1.13)	(0.83)	(0.73)	(0.66)	(0.55)	(0.47)	(0.44)	(0.42)
$\mathbb{E}\left[g^2\right]$	15.41	16.45	17.37	18.25	18.88	19.53	20.07	20.63	21.11	21.59	22.06	22.50
	(1.07)	(0.95)	(0.87)	(0.88)	(0.89)	(0.91)	(0.89)	(0.86)	(0.84)	(0.82)	(0.80)	(0.78)
$\mathbb{E}^{\mathbb{Q}}\left[r^2\right]$	45.30	45.55	46.47	47.07	47.39	47.86	48.40	48.81	49.14	49.33	49.40	49.94
	(3.53)	(3.03)	(2.92)	(2.66)	(2.43)	(2.38)	(2.34)	(2.32)	(2.29)	(2.26)	(2.19)	(2.18)
$\mathbb{E}^{\mathbb{Q}}\left[l^2\right]$	30.46	31.16	32.01	32.50	32.71	33.03	33.39	33.65	33.81	33.83	33.74	34.09
	(2.58)	(2.26)	(2.24)	(2.04)	(1.85)	(1.84)	(1.84)	(1.84)	(1.85)	(1.83)	(1.77)	(1.78)
$\mathbb{E}^{\mathbb{Q}}\left[g^2\right]$	14.84	14.39	14.46	14.57	14.68	14.84	15.01	15.16	15.33	15.50	15.66	15.85
	(0.98)	(0.80)	(0.73)	(0.68)	(0.64)	(0.61)	(0.59)	(0.58)	(0.57)	(0.56)	(0.56)	(0.55)
QRP	19.04	19.07	19.07	18.89	19.43	20.01	20.44	20.77	21.10	21.19	21.04	21.32
	(1.35)	(1.43)	(1.43)	(1.36)	(1.41)	(1.44)	(1.41)	(1.43)	(1.46)	(1.45)	(1.43)	(1.44)
QRP^l	19.61	21.13	21.98	22.57	23.64	24.70	25.50	26.23	26.89	27.29	27.43	27.97
	(1.42)	(1.54)	(1.63)	(1.65)	(1.69)	(1.71)	(1.69)	(1.71)	(1.74)	(1.73)	(1.69)	(1.70)
QRP^g	0.56	2.05	2.91	3.68	4.20	4.69	5.06	5.46	5.78	6.10	6.40	6.65
	(0.33)	(0.37)	(0.40)	(0.46)	(0.44)	(0.46)	(0.47)	(0.47)	(0.47)	(0.47)	(0.47)	(0.48)

Table 2: **Principal Component Analysis**

In this table, we report in percentage the explanatory power of each of the first three principal components, and their total explanatory power, for a number of different information sets. These include the term structure of the loss and gain components of the physical expected quadratic payoff, the risk-neutral expected quadratic payoff, and the quadratic risk premium, each separately. We also report the explanatory power of the first three principal components of the term structure of components of the physical expected quadratic payoff together with the term structure of the components of the quadratic risk premium. The sample period is from January 1996 to December 2015.

Principal Component	1	2	3	First 3
Information Sets	Explanatory Power			
$\mathbb{E}\left[l^2\right]$ and $\mathbb{E}\left[g^2\right]$	58.01	37.10	3.13	98.24
$\mathbb{E}^{\mathbb{Q}}\left[l^2\right]$ and $\mathbb{E}^{\mathbb{Q}}\left[g^2\right]$	87.33	8.78	2.76	98.87
QRP^l and QRP^g	73.05	15.22	6.18	94.45
$\mathbb{E}\left[l^2\right]$, $\mathbb{E}\left[g^2\right]$, QRP^l and QRP^g	56.73	26.73	7.93	91.39

Table 3: **Risk-Neutral Moments RMSEs**

In this table we report the root mean squared error

$$RMSE \equiv \sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(Mom_t^{Mkt} - Mom_t^{Mod}\right)^2},$$

where Mom_t^{Mkt} is the time t risk-neutral moment value observed on the market, Mom_t^{Mod} is the corresponding model-implied equivalent. All variance RMSEs are in annual percentage units. The sample period is from January 1996 to December 2015.

τ	Quadratic Loss					Quadratic Gain				
	AFT0	**AFT1**	**AFT2**	**AFT3**	**AFT4**	**AFT0**	**AFT1**	**AFT2**	**AFT3**	**AFT4**
2	1.08	0.97	1.14	0.77	0.68	2.56	0.69	1.50	1.65	0.82
3	1.47	0.60	0.81	0.62	0.50	2.12	0.72	1.13	1.37	0.63
4	1.73	0.40	0.57	0.52	0.44	1.99	0.78	0.95	1.14	0.69
5	1.82	0.45	0.46	0.49	0.46	1.93	0.85	0.91	1.08	0.80
6	1.77	0.66	0.55	0.64	0.54	1.98	0.92	0.98	1.06	1.00
7	1.74	0.88	0.74	0.80	0.66	2.07	0.99	1.03	1.07	1.19
8	1.71	1.06	0.89	0.97	0.83	2.15	1.06	1.07	1.08	1.37
9	1.70	1.23	1.05	1.12	0.95	2.21	1.12	1.10	1.03	1.51
10	1.86	1.36	1.20	1.23	1.03	2.28	1.19	1.13	1.00	1.63
11	1.99	1.49	1.29	1.32	1.11	2.38	1.21	1.12	1.00	1.72
12	2.15	1.60	1.39	1.36	1.22	2.45	1.29	1.06	1.07	1.80
Avg	1.73	0.97	0.92	0.89	0.77	2.19	0.98	1.09	1.14	1.20
τ	Volatility					Skewness				
	AFT0	**AFT1**	**AFT2**	**AFT3**	**AFT4**	**AFT0**	**AFT1**	**AFT2**	**AFT3**	**AFT4**
2	0.74	1.09	1.04	0.69	0.31	0.97	0.19	0.26	0.30	0.47
3	0.33	0.70	0.70	0.61	0.11	0.95	0.16	0.19	0.22	0.27
4	0.51	0.44	0.43	0.47	0.18	0.97	0.15	0.15	0.17	0.19
5	0.47	0.29	0.24	0.31	0.21	0.92	0.13	0.12	0.13	0.13
6	0.43	0.21	0.19	0.19	0.21	0.88	0.11	0.09	0.09	0.09
7	0.36	0.29	0.26	0.23	0.18	0.85	0.08	0.09	0.08	0.09
8	0.22	0.40	0.39	0.33	0.15	0.81	0.08	0.11	0.09	0.11
9	0.20	0.55	0.52	0.44	0.12	0.79	0.09	0.13	0.09	0.15
10	0.36	0.67	0.64	0.50	0.16	0.75	0.11	0.17	0.10	0.17
11	0.54	0.81	0.78	0.61	0.20	0.72	0.13	0.19	0.12	0.19
12	0.71	0.93	0.93	0.75	0.28	0.74	0.14	0.19	0.15	0.22
Avg	0.44	0.58	0.56	0.47	0.19	0.85	0.12	0.15	0.14	0.19

Table 4: **Quadratic Risk Premium RMSEs**

In this table we report the root mean squared error

$$RMSE \equiv \sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(QRP_t^{Mkt} - QRP_t^{Mod}\right)^2},$$

where QRP_t^{Mkt} is the time t Quadratic Risk Premium value observed on the market, QRP_t^{Mod} is the corresponding model-implied equivalent. All variance RMSEs are in annual percentage units. The sample period is from January 1996 to December 2015.

τ	Loss QRP					Gain QRP				
	RND0	**RND1**	**RND2**	**RND3**	**RND4**	**RND0**	**RND1**	**RND2**	**RND3**	**RND4**
3	9.86	6.39	3.92	3.55	6.36	1.72	1.95	2.79	2.80	1.12
6	11.53	9.21	4.08	6.04	2.73	1.53	1.62	0.93	1.41	0.92
9	12.53	10.82	3.85	8.01	2.73	1.69	1.70	1.18	1.43	1.23
12	13.21	11.63	4.09	9.16	2.55	1.95	1.93	1.62	1.73	1.52
Avg	11.78	9.51	3.99	6.69	3.59	1.72	1.80	1.63	1.84	1.20
τ	Net QRP (i.e Loss QRP - Gain QRP)					Skewness RP (Loss QRP+Gain QRP)				
	RND0	**RND1**	**RND2**	**RND3**	**RND4**	**RND0**	**RND1**	**RND2**	**RND3**	**RND4**
3	7.33	4.86	2.63	3.28	2.77	10.90	7.31	4.08	4.85	6.60
6	7.85	6.15	2.16	3.90	1.17	11.98	9.84	3.71	6.59	2.27
9	7.42	6.11	1.70	4.12	0.94	11.88	10.25	3.01	7.52	2.13
12	7.05	5.88	1.32	4.12	0.91	11.80	10.40	2.72	8.06	1.92
Avg	7.41	5.75	1.95	3.86	1.45	11.64	9.45	3.38	6.76	3.23

Figure 1: Term Structure of Expected Quadratic Payoff

In this figure, in Panel A we plot the mean S&P 500 risk-neutral expected quadratic payoff and its loss and gain components for maturities of 1, 3, 6, 9 and 12 months. In Panel B, we plot the same quantities for the physical excepted quadratic payoff also for maturities of 1, 3, 6, 9 and 12 months. The sample period is from January 1996 to December 2015.

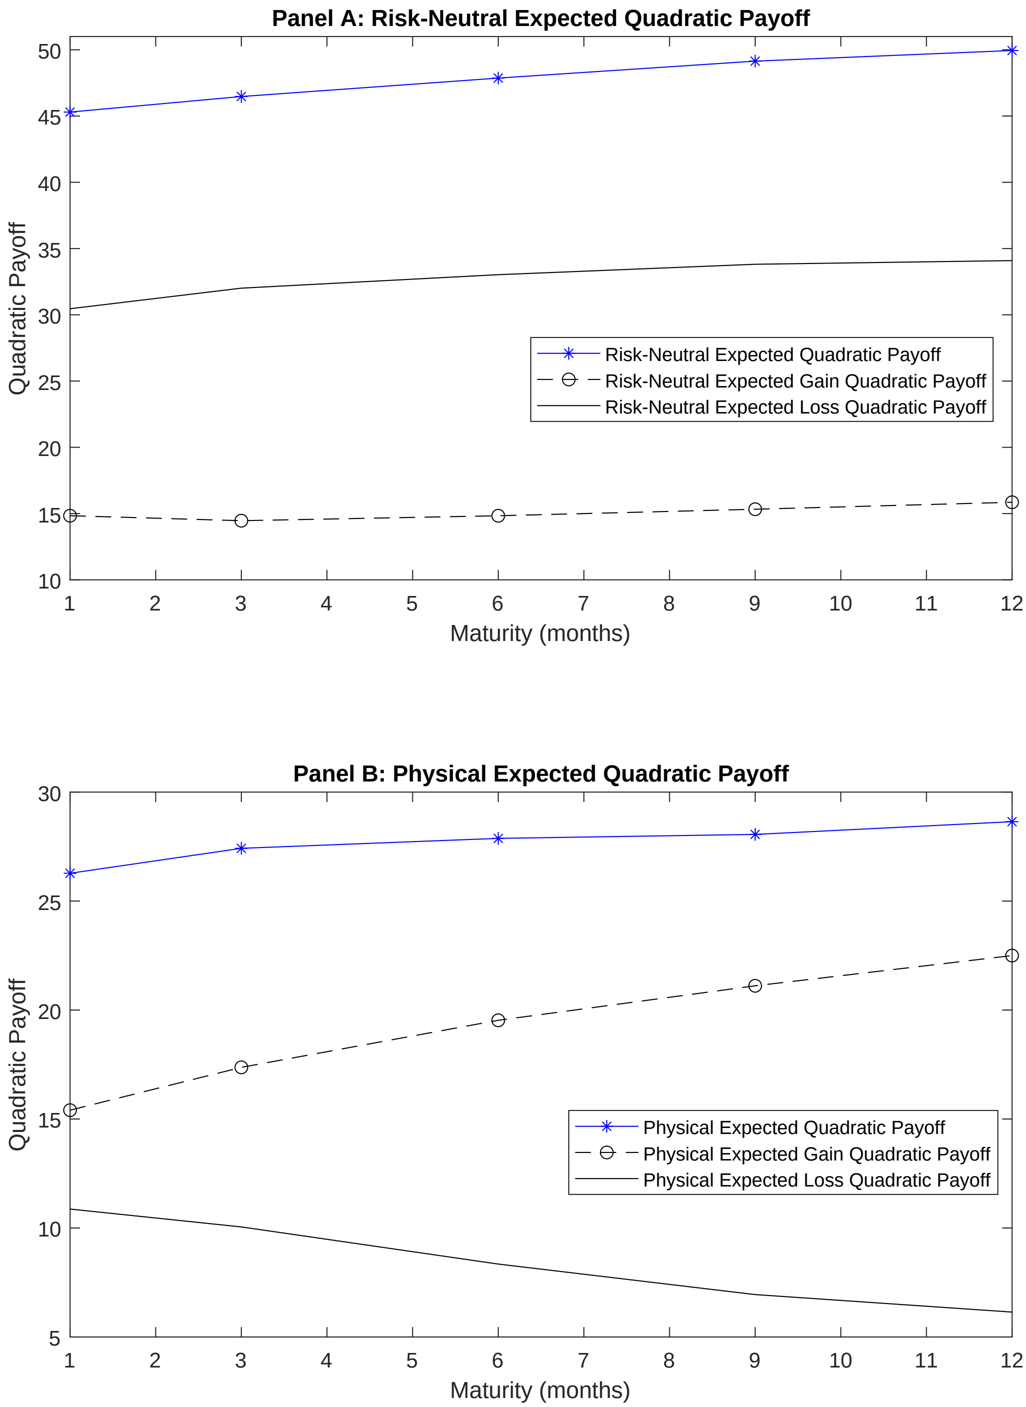

Figure 2: Level and Slope of Expected Quadratic Payoff Term Structure

In this figure, in Panel A we plot the level (6-month maturity) of the S&P 500 risk-neutral ($\mathbb{Q}$) expected quadratic payoff and its loss and gain components. In Panel B, we plot the slope (12-month minus the 2-month maturity) of the S&P 500 risk-neutral expected quadratic payoff and its loss and gain components. In Panel C and D, we plot the same quantities for the physical ($\mathbb{P}$) excepted quadratic payoff also for the level and slope, respectively. The sample period is from January 1996 to December 2015.

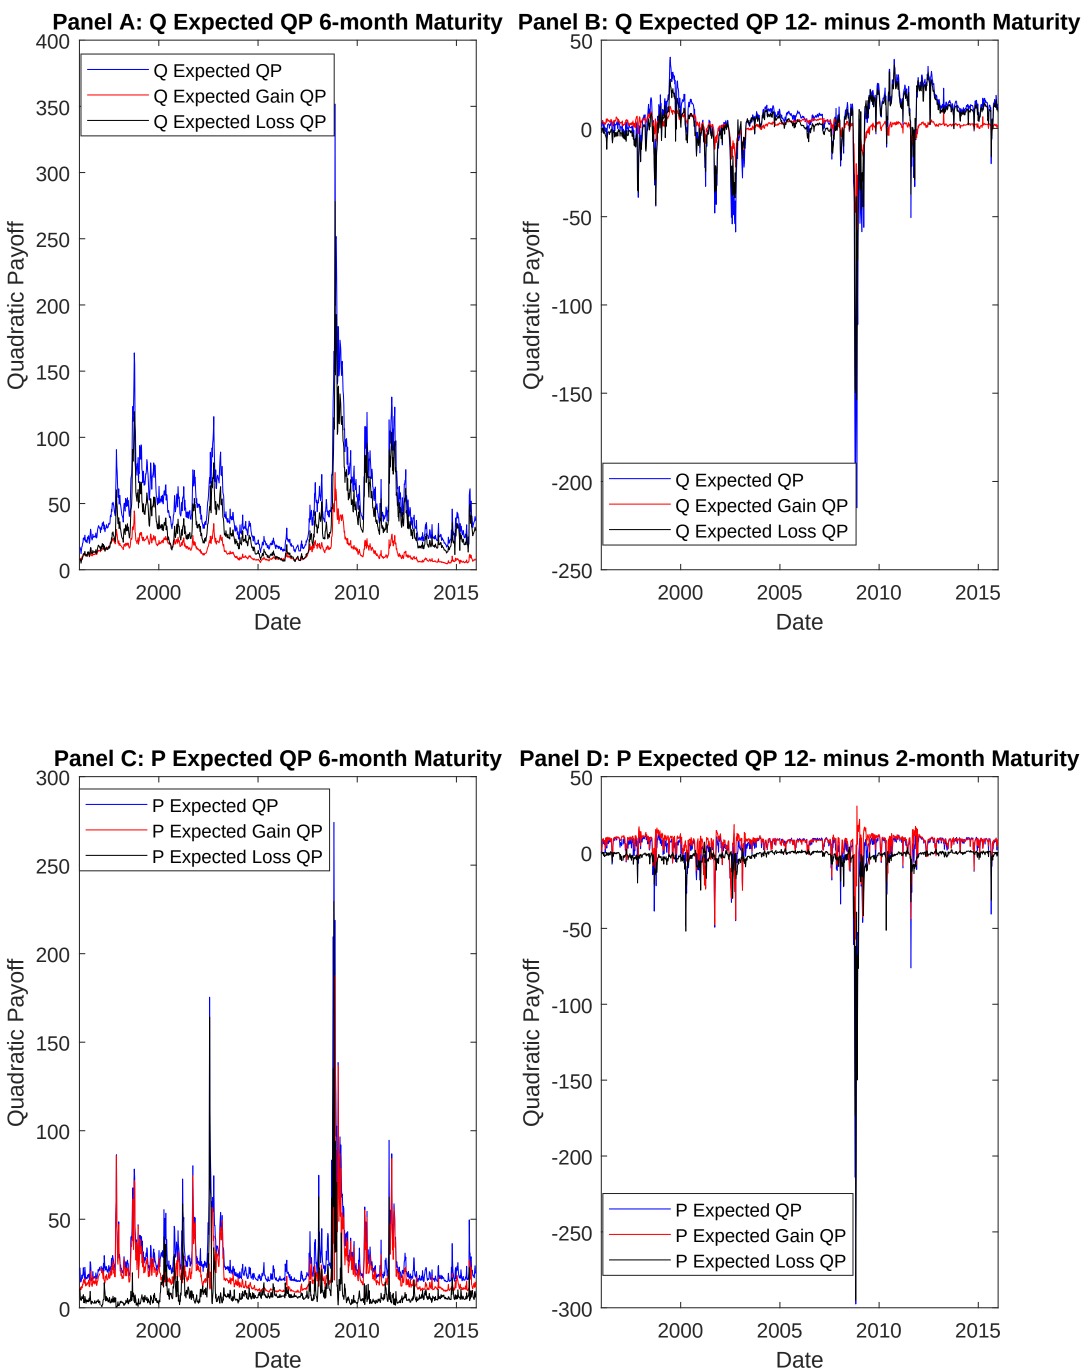

Figure 3: Observed and models implied average term structure of risk-neutral moments and risk-premium

In this figure we plot the observed and models implied average term structure of risk-neutral moments (the first two rows) and quadratic risk premium (the third row). For the risk premium, we use the most flexible specification of the AFT model (AFT4). The sample period is from January 1996 to December 2015.