

The Term Structures of Expected Loss and Gain Uncertainty

by Bruno Feunou,¹ Ricardo Lopez Aliouchkin,² Roméo Tédongap³
and Lai Xu⁴

¹ Financial Markets Department
Bank of Canada, Ottawa, Ontario, Canada K1A 0G9
bfeunou@bankofcanada.ca

² Syracuse University

³ ESSEC Business School

⁴ Syracuse University



Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

Acknowledgements

We would like to thank the editor Francis X. Diebold and two anonymous referees for their helpful comments, which greatly improved the paper. Feunou gratefully acknowledges financial support from the Canadian Derivatives Institute. Lopez Aliouchkin and Tédongap acknowledge the research grant from the Thule Foundation's Skandia research programs on "Long-Term Savings." The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of Canada.

Abstract

We document that the term structures of risk-neutral expected loss and gain uncertainty on S&P 500 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.'s (2015) 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: (i) the inclusion of jumps, (ii) disentangling the price of negative jump risk from its positive analog in the stochastic discount factor specification, and (iii) specifying three latent factors.

Topics: Asset pricing; Econometric and statistical methods

JEL code: G12

1 Introduction

Financial economists have long agreed that to better understand asset returns, and also uncertainty about these returns, it is 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 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 also leads 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, is among 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 understanding market views about short- and long-term loss and gain potential. When dealing with expectations of asset returns uncertainty, it is also important to distinguish between physical and risk-neutral expectations. On the one hand, physical expectations of uncertainty measure the degree to which investors anticipate that they could be wrong about their returns forecast. On the other hand, risk-neutral expectations of uncertainty additionally indicate how much investors are willing to pay for risk hedging or would require for risk compensation. The shape of the term structure of risk-neutral expectations of loss (gain) uncertainty reflects both the expected path of future loss (gain) volatility and different risk premia associated with downside (upside) risk at different maturities.

The primary goal of this article is to provide 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. We assume we observe returns at regular intra-month time

intervals of length δ . The monthly realized return $r_{t-1,t}$ and the monthly realized variance $RV_{t-1,t}$ are defined by aggregating $r_{t-1+j\delta}$ and $r_{t-1+j\delta}^2$, respectively:

$$r_{t-1,t} = \sum_{j=1}^{1/\delta} r_{t-1+j\delta} \quad \text{and} \quad RV_{t-1,t} = \sum_{j=1}^{1/\delta} r_{t-1+j\delta}^2, \quad (1)$$

where $1/\delta$ is the number of higher-frequency returns in a monthly period (e.g., $\delta = 1/21$ for daily returns) and $r_{t-1+j/21}$ denotes the j th daily return of the monthly period starting from day $t-1$ and ending on day t . The loss component of realized variance cumulates higher-frequency squared losses, $l_{t-1+j\delta}$, while the gain component sums up higher-frequency squared gains, $g_{t-1+j\delta}$, that is,

$$RV_{t-1,t}^l = \sum_{j=1}^{1/\delta} l_{t-1+j\delta}^2 \quad \text{and} \quad RV_{t-1,t}^g = \sum_{j=1}^{1/\delta} g_{t-1+j\delta}^2. \quad (2)$$

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 the loss and gain realized variance is not feasible, nor are loss and gain variance swaps traded such that their strikes could then be observed measures for these risk-neutral expectations. To illustrate this difficulty, consider the risk-neutral expectation of the monthly realized variance based on daily returns. It is the sum of risk-neutral expectations of daily squared returns, $\mathbb{E}^{\mathbb{Q}} [r_{t-1+j\delta}^2]$, each of which would be computable in theory, for example, using the formula of Bakshi et al. (2003) for the price of the quadratic contract. Empirically, this would require data on one-day-to-maturity out-of-the-money options, which are not currently available. The same applies to risk-neutral expectations of the monthly loss and gain realized variance. 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. The quadratic payoff is one of such measures and is the focus in this article.

The quadratic payoff is the square of the realized return over the investment horizon; that is, the monthly quadratic payoff is simply defined as the squared monthly realized return, $r_{t-1,t}^2$. The quadratic payoff and the realized variance are therefore related as follows:

$$r_{t-1,t}^2 = RV_{t-1,t} + 2RA_{t-1,t}, \quad \text{where} \quad RA_{t-1,t} = \sum_{i=1}^{1/\delta-1} \sum_{j=1}^{1/\delta-i} r_{t-1+j\delta} r_{t-1+j\delta+i\delta}, \quad (3)$$

and $RA_{t-1,t}$ can be defined as the realized autocovariance. Like the realized variance, the quadratic payoff is also a model-free 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. Formally and in monthly terms, this means $l_{t-1,t}^2$ and $g_{t-1,t}^2$, respectively. Also similar to the realized variance, the quadratic payoff is the sum of its loss and gain components. Contrary to the realized variance, both physical and risk-neutral expectations of quadratic loss and gain for various horizons can be consistently estimated or measured from the data. We provide full details in Internet Appendix Section A.1. We therefore rely on the quadratic payoff as the measure of asset returns uncertainty when analyzing the term structure of expected loss uncertainty and gain uncertainty.

Using a large panel of S&P 500 Index options data with time to maturity ranging from 1 month to 12 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&P 500 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 (with a slope of -4.73 percentage square units), while the average term structure of the risk-neutral expected quadratic loss is upward sloping (with a slope of 3.63 percentage square units). 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 of stocks is more expensive than hedging the short-term loss potential. Second, the average term structure of the physical expected quadratic gain is upward sloping (with a slope of 7.09 percentage square units), while the average term structure of the risk-neutral expected quadratic gain is slightly upward sloping and almost flat (with a slope of 1.01 percentage square units). 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 of stocks is almost as costly as speculating on the long-term gain potential.

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 risk-neutral expectations. 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 the 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 find that both term structures are upward sloping (with slopes of 8.36 and 6.09 percentage square units, 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 that 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 the 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 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

¹Some 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.

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 dynamics 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, e.g., Eraker, 2004, Santa-Clara and Yan, 2010, Christoffersen et al., 2012, and Bates, 2012) and highlights its inability to account for the actual joint 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) investigate the ability of existing structural models to fit 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 theoretical arguments supporting the splitting of the total realized variance into 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 details on the estimation of the AFT model. Section 5 evaluates the ability of the AFT model and its variants to fit 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 $r_{t,t+\tau} = \ln(S_{t+\tau}/S_t)$. Both the log return $r_{t,t+\tau}$ and the quadratic payoff $r_{t,t+\tau}^2$ are subject to a gain-loss decomposition as follows:

$$r_{t,t+\tau} = g_{t,t+\tau} - l_{t,t+\tau} \quad \text{and} \quad r_{t,t+\tau}^2 = g_{t,t+\tau}^2 + l_{t,t+\tau}^2, \quad (4)$$

where the gain $g_{t,t+\tau} = \max(0, r_{t,t+\tau})$ and the loss $l_{t,t+\tau} = \max(0, -r_{t,t+\tau})$ represent the positive and negative parts of the asset payoff, respectively. In other words, the gain and 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 in 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 $\mathbb{E}_t^{\mathbb{Q}}[l_{t,t+\tau}^2]$ and $\mathbb{E}_t^{\mathbb{Q}}[g_{t,t+\tau}^2]$, and of the physical expectations $\mathbb{E}_t^{\mathbb{P}}[l_{t,t+\tau}^2]$ and $\mathbb{E}_t^{\mathbb{P}}[g_{t,t+\tau}^2]$, 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 risk or would 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 physical expectations of the quadratic payoff as the quadratic risk premium (QRP), where 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

$$QRP_t^l(\tau) \equiv \mathbb{E}_t^{\mathbb{Q}}[l_{t,t+\tau}^2] - \mathbb{E}_t^{\mathbb{P}}[l_{t,t+\tau}^2] \quad \text{and} \quad QRP_t^g(\tau) \equiv \mathbb{E}_t^{\mathbb{P}}[g_{t,t+\tau}^2] - \mathbb{E}_t^{\mathbb{Q}}[g_{t,t+\tau}^2]. \quad (5)$$

Eq. (5) shows that the 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 physical expected quadratic loss and gain naturally leads to examining the term structures of the loss and gain QRPs.

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

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

$$\mu_n^{\mathbb{Q}+}(t, \tau) \equiv \mathbb{E}_t^{\mathbb{Q}}[g_{t,t+\tau}^n], \quad \mu_n^{\mathbb{Q}-}(t, \tau) \equiv \mathbb{E}_t^{\mathbb{Q}}[l_{t,t+\tau}^n], \quad \text{and} \quad \mu_n^{\mathbb{Q}}(t, \tau) \equiv \mathbb{E}_t^{\mathbb{Q}}[r_{t,t+\tau}^n], \quad (6)$$

$$\mu_n^{\mathbb{P}+}(t, \tau) \equiv \mathbb{E}_t^{\mathbb{P}}[g_{t,t+\tau}^n], \quad \mu_n^{\mathbb{P}-}(t, \tau) \equiv \mathbb{E}_t^{\mathbb{P}}[l_{t,t+\tau}^n], \quad \text{and} \quad \mu_n^{\mathbb{P}}(t, \tau) \equiv \mathbb{E}_t^{\mathbb{P}}[r_{t,t+\tau}^n]. \quad (7)$$

To understand the difference between $\mu_2^{\mathbb{Q}+}(t, \tau)$ and $\mu_2^{\mathbb{Q}-}(t, \tau)$, we follow Proposition 2 of Duffie et al. (2000),

$$\mu_2^{\mathbb{Q}-}(t, \tau) = \frac{\mathbb{E}_t^{\mathbb{Q}}[r_{t,t+\tau}^2] + \Lambda^{\mathbb{Q}}(t, \tau)}{2}, \quad \mu_2^{\mathbb{Q}+}(t, \tau) = \frac{\mathbb{E}_t^{\mathbb{Q}}[r_{t,t+\tau}^2] - \Lambda^{\mathbb{Q}}(t, \tau)}{2}, \quad (8)$$

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

$$\Lambda^{\mathbb{Q}}(t, \tau) = \frac{2}{\pi} \int_0^{+\infty} \frac{\text{Im} \left(\varphi_{t,\tau}^{(2)}(-iv) \right)}{v} dv, \quad (9)$$

with $\varphi_{t,\tau}(\cdot)$ being the time- t conditional risk-neutral moment-generating function of $r_{t,t+\tau}$ and $\varphi_{t,\tau}^{(2)}(\cdot)$ its second-order derivative, and $\text{Im}(\cdot)$ refers to the imaginary coefficient of a complex num-

ber. From Eq. (8), 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 $\mathbb{E}_t^{\mathbb{Q}}[r_{t,t+\tau}^2]$ and the term structure of $\Lambda^{\mathbb{Q}}(t, \tau)$. Several papers in the literature have already dealt successfully with $\mathbb{E}_t^{\mathbb{Q}}[r_{t,t+\tau}^2]$, and the consensus seems to be that a two-factor diffusion model provides 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

$$\Lambda^{\mathbb{Q}}(t, \tau) = \lim_{\bar{v} \rightarrow \infty} \left\{ \sum_{j=1}^{\infty} \frac{(-1)^j \bar{v}^{2j-1}}{(2j-1)(2j-1)!} \mu_{2j+1}^{\mathbb{Q}}(t, \tau) \right\}, \quad (10)$$

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 follows:

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

where $\kappa_n^{\mathbb{Q}}(t, \tau) \equiv \mathbb{E}_t^{\mathbb{Q}} \left[\left(r_{t,t+\tau} - \mu_1^{\mathbb{Q}}(t, \tau) \right)^n \right]$. 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 exists and is 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:

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

where S_t is the time- t price of the 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

$$\mathbb{E}_t^{\mathbb{Q}}[r_{t,t+\tau}^2] = e^{r_f \tau} V_t(\tau), \quad (12)$$

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 Eq. (11) by setting up a total of 1,001 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 1,001 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 (see, e.g., Conrad et al., 2013, and others). Lastly, 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 $\tau = 30, 60, \dots, 360$ days.

²Since 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 Eq. (11) is the sum of a portfolio of OTM call options and a portfolio of OTM put options:

$$V_t(\tau) = V_t^g(\tau) + V_t^l(\tau), \quad (13)$$

where

$$V_t^l(\tau) = \int_0^{S_t} \frac{1 + \ln(S_t/K)}{K^2/2} P_t(\tau; K) dK \quad \text{and} \quad V_t^g(\tau) = \int_{S_t}^{\infty} \frac{1 - \ln(K/S_t)}{K^2/2} C_t(\tau; K) dK. \quad (14)$$

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 Section A.6 of the Internet Appendix accompanying this paper. Hence, the risk-neutral expectations of quadratic loss and gain are

$$\mathbb{E}_t^{\mathbb{Q}}[l_{t,t+\tau}^2] = e^{r_f \tau} V_t^l(\tau) \quad \text{and} \quad \mathbb{E}_t^{\mathbb{Q}}[g_{t,t+\tau}^2] = e^{r_f \tau} V_t^g(\tau). \quad (15)$$

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 then have

$$\mathbb{E}_t[r_{t,t+\tau}^2] = \mu_{t,\tau}^2 + \sigma_{t,\tau}^2 \quad \text{and} \quad \begin{cases} \mathbb{E}_t[l_{t,t+\tau}^2] &= (\mu_{t,\tau}^2 + \sigma_{t,\tau}^2) \Phi\left(-\frac{\mu_{t,\tau}}{\sigma_{t,\tau}}\right) - \mu_{t,\tau} \sigma_{t,\tau} \phi\left(\frac{\mu_{t,\tau}}{\sigma_{t,\tau}}\right) \\ \mathbb{E}_t[g_{t,t+\tau}^2] &= (\mu_{t,\tau}^2 + \sigma_{t,\tau}^2) \Phi\left(\frac{\mu_{t,\tau}}{\sigma_{t,\tau}}\right) + \mu_{t,\tau} \sigma_{t,\tau} \phi\left(\frac{\mu_{t,\tau}}{\sigma_{t,\tau}}\right), \end{cases} \quad (16)$$

under the log-normality assumption, where $\Phi(\cdot)$ and $\phi(\cdot)$ are the standard normal cumulative distribution functions and density, respectively.

The first part of (16) implies that the difference between the physical expected quadratic payoff $\mathbb{E}_t[r_{t,t+\tau}^2]$ and the physical expected realized variance $\sigma_{t,\tau}^2$ is due to the non-zero drift ($\mu_{t,\tau} \neq 0$). If the drift equals zero, the expected quadratic payoff is exactly the same as the expected realized variance. In Figure A7 of the Internet Appendix, we plot the average term structure of the drift ($\mu_{t,\tau}$) and the squared drift ($\mu_{t,\tau}^2$). As expected, the term structure of the drift is essentially flat. However, the term structure of the squared drift is increasing in the investment horizon,

which explains the increasing discrepancy between the expected quadratic payoff and the expected realized variance at longer horizons. We also report descriptive statistics of both quantities in Table A4. We find that both the drift and the squared drift of any investment horizon exhibit significant variations over time. Overall, these results highlight the importance of studying the term structure of the quadratic payoff.

The second part of Eq. (16) implies that the wedge between the physical expected quadratic gain and loss is a function of the drift $\mu_{t,\tau}$. If the drift equals zero, then estimates of the physical expected quadratic gain and loss are equal. This is clearly a consequence of the normality assumption. In general, the wedge between the physical expected quadratic gain and loss could also be driven by other factors, for example, skewness. In the Internet Appendix (Section A.10), we relax the normality assumption and instead assume that log returns $r_{t,t+\tau}$ follow a binormal distribution. This is an analytically tractable distribution that accommodates empirically plausible values of skewness and kurtosis, and nests the familiar Gaussian distribution (for details, see Feunou et al., 2013). In Table A5 of the Internet Appendix, we report the correlations between the physical expected quadratic loss (gain) estimated under the assumption of normal and binormal distributed log returns for a 1- to 12-month investment horizon. We find that correlations between these two physical expected quadratic loss (gain) quantities range between 0.984 (0.980) and 0.996 (0.993). Further, in Table A6 of the Internet Appendix, we find that the average term structures of these quantities have consistent patterns and similar values. Taken together, this evidence suggests that our main results are robust to the distributional assumption of returns.

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 Eq. (16) 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 dominate 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:

$$\begin{aligned}
RV_{t,t+\tau} = & c + \alpha VIX_t^2 + \beta^m C_{t-21,t} + \beta^w C_{t-5,t} + \beta^d C_t \\
& + \gamma^m J_{t-21,t} + \gamma^w J_{t-5,t} + \gamma^d J_t \\
& + \delta^m l_{t-21,t} + \delta^w l_{t-5,t} + \delta^d l_{t-1,t} + \varepsilon_{t+\tau}^{(\tau)},
\end{aligned} \tag{17}$$

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 a 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 (17), for values of $\tau = 21, 42, \dots, 252$ days. Likewise, the conditional mean $\mu_{t,\tau}$ is the fitted time series from the regression (17) 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 (see, e.g., Carr and Wu, 2009). Following Bakshi et al. (2003), we restrict the sample to OTM 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, for example, call options with a price higher than the underlying price and options with less than seven 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 j th 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 variances 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 $\mathbb{E}_t^{\mathbb{Q}} [l_{t,t+\tau}^2]$ or $\mathbb{E}_t^{\mathbb{Q}} [g_{t,t+\tau}^2]$ for each maturity τ , we use options with maturity close to τ and do interpolations.⁴ In the top left panel of Figure 1, we plot the time-series average of the 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 the risk-neutral expected gain quadratic payoff is flat.

To investigate time variations in these term structures, in the two middle panels of Figure 1, we plot the 6-month maturity (the level) and the 12- minus 2-month maturity (the slope) for the risk-neutral expected quadratic 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 in the short run.

Relationship to Dew-Becker et al. (2017) Dew-Becker et al. compute the term structure of forward variance prices as $F_t^{rv,\tau} \equiv \mathbb{E}_t^{\mathbb{Q}} [RV_{t+\tau-1,t+\tau}]$. The forward variance price $F_t^{rv,\tau}$ is essentially the month- t risk-neutral expectation of the realized variance from end of month $t + \tau - 1$ to end of month $t + \tau$. The authors compute these forward prices using traded variance swaps. Since

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

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

traded loss and gain variance swaps do not exist, one cannot estimate risk-neutral expectations of the loss and gain components using their approach. Previous literature (see, e.g., Kilic and Shaliastovich, 2019) has shown that loss and gain components are important for asset prices. Our methodology allows us to compute the term structure of forward prices not only for the quadratic payoff but also for its loss and gain components as $F_t^{r,\tau} \equiv \mathbb{E}_t^{\mathbb{Q}} [r_{t+\tau-1,t+\tau}^2]$, $F_t^{l,\tau} \equiv \mathbb{E}_t^{\mathbb{Q}} [l_{t+\tau-1,t+\tau}^2]$, and $F_t^{g,\tau} \equiv \mathbb{E}_t^{\mathbb{Q}} [g_{t+\tau-1,t+\tau}^2]$, respectively.

In Panel A of Figure A1 in the Internet Appendix, we plot the term structure of average forward prices for the risk-neutral expected quadratic payoff and its loss and gain components. We find that the estimated term structure of forward prices for the quadratic payoff is concave, with both the level and slope similar to the forward prices in Dew-Becker et al. (2017). This finding is despite the differences in sample periods and the fact that our forward prices are computed for the quadratic payoff while their forward prices are computed from traded variance swaps. Most importantly, as mentioned above, in contrast to Dew-Becker et al. (2017) we are able to separately estimate the forward prices for the quadratic loss and gain, while loss and gain variance swaps do not exist. This approach allows us to investigate not only the term structure of the quadratic payoff but also its components. We see that the term structures of the average forward prices for the quadratic loss and gain are in general upward sloping. We also find that, across all horizons, the quadratic loss forward prices are higher than the quadratic gain forward prices.

2.5.2 *The Term Structure of the Physical Expected Quadratic Payoff*

In the top right panel of Figure 1, we plot the time-series average of the term structure of the 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 finding suggests that investors face more uncertainty about losses in the short run relative to 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 finding suggests that investors face more uncertainty about gains in the long run relative to the 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 the two bottom panels of Figure 1, 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 in the short run.

Further, in the top two panels of 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 relative to 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 the time-series means of the risk-neutral and physical expected quadratic payoff together with their loss and gain components. For each 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 mean values for the risk-neutral expected quadratic loss are much higher than the values for 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.

Rolling Window Parameter Estimation We use the full sample to estimate the drifts in returns and the expectations of realized variances over different periods. These estimated quantities enter Eq. (16) to compute the expectations of the truncated returns. Such an approach may incur forward-looking bias in the parameter estimation stage. To check the robustness of our results, we estimate drifts and expected realized variances using a set of different rolling windows: 60, 72, 84, 96, 108, and 120 months of daily data. Detailed comparisons to the full-sample expected quadratic gain and loss are discussed in the Internet Appendix. We plot the correlation between the rolling-window and full-sample expected quadratic gain in Figure A9 of the Internet Appendix, which reveals that all the correlations are well above 0.7. Except for the expected quadratic gain using a 60-month window size, all other correlations are higher than 0.77. The correlation can be as high as 0.88 for the expected quadratic loss based on a 120-month rolling window.

To further compare the full-sample and rolling-window expected truncated returns, we fix the rolling window size to 120 months, and for each investment horizon τ , we estimate time-series regressions of the following form:

$$\begin{aligned}\mathbb{E}[g_{FS,\tau}^2] &= \alpha + \beta_{roll}^g \mathbb{E}[g_{roll,\tau}^2] + \varepsilon \\ \mathbb{E}[l_{FS,\tau}^2] &= \alpha + \beta_{roll}^l \mathbb{E}[l_{roll,\tau}^2] + \varepsilon.\end{aligned}\tag{18}$$

If these expectations constructed from rolling-window parameters are exactly the same as the ones constructed from full-sample parameters, we should find that β_{roll}^g and β_{roll}^l are not statistically significantly different from one. In Figure A10 of the Internet Appendix, we plot these coefficients for investment horizons from 1 to 12 months. Panel A shows that β_{roll}^g is not statistically different from one over horizons of 6, 7, and 8 months. On the other hand, we find in Panel B that β_{roll}^l is not statistically different from one over a 3-month horizon. Overall, because of the difficulty predicting returns over short horizons, the quantities constructed from the full-sample parameters inevitably differ from the rolling-window counterparts.

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 the 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, and 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 the 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 represents the insurance cost (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 three principal components. We find that the first three principal components are enough to explain 91.39% of the variation in 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 56.73%, the second explains 26.73%, and the third explains 7.93% 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 term 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 for three reasons. First, it is built to disentangle the dynamics of positive and negative jumps. Second, it is a three-factor framework, which would maximize the model's chances of fitting the term structure of the expected quadratic loss and gain and their risk premia since we find that three principal components are needed to fit the targeted term structures in Section 2.5.4. Third, 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 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}):

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

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

for the jump measure μ , and is assumed to be

$$\nu_t^{\mathbb{Q}}(dx, dy) = \left\{ \left(c_t^- 1_{\{x < 0\}} \lambda_- e^{-\lambda_- |x|} + c_t^+ 1_{\{x > 0\}} \lambda_+ e^{-\lambda_+ |x|} \right) 1_{\{y=0\}} + c_t^- 1_{\{x=0, y < 0\}} \lambda_- e^{-\lambda_- |y|} \right\} dx \otimes dy, \quad (19)$$

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

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

These three factors have distinctive features: V_{2t} is a pure-diffusion process, V_{3t} is a pure-jump process, and innovation in V_{1t} combines 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.

To better understand the ability of 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. The first is the number of factors. Compared to the three-factor framework, we ask whether two factors are enough and which two-factor alternatives generate the best fit. The second dimension is the model's ability to differentiate between the negative and positive jump distribution. We ask whether 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 are no jumps. This corresponds to the two-factor diffusion model studied extensively in Christoffersen et al. (2009). This is equivalent to suppressing all the jump 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, a subject of much debate in the option pricing literature.
- AFT3: The expected negative jump variation equals the expected positive jump variation. This corresponds to a three-factor model that 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 the existing option pricing and variance swap models as it does not differentiate between the intensity of positive and negative jumps (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 three-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.⁵ In the Internet Appendix, we compare two three-factor models in which $\eta = 0$ and $\eta \neq 0$ 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 the statistical properties of the 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 that are 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 for 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

⁵Several papers, 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&P 500 options.

risk-neutral dynamic. In particular, the price of jump risk, $\Psi_t(x, y)$, is given by

$$\Psi_t(x, y) \equiv \frac{\nu_t^{\mathbb{Q}}(dx, dy)}{\nu_t^{\mathbb{P}}(dx, dy)} - 1 \quad \text{where} \quad \frac{\nu_t^{\mathbb{Q}}(dx, dy)}{\nu_t^{\mathbb{P}}(dx, dy)} = \begin{cases} \frac{c_t^+}{c_t^{\mathbb{P}+}} \frac{\lambda_+}{\lambda_+^{\mathbb{P}}} \exp(-(\lambda_+ - \lambda_+^{\mathbb{P}})x) & x > 0 \quad y = 0 \\ \frac{c_t^-}{c_t^{\mathbb{P}-}} \frac{\lambda_-}{\lambda_-^{\mathbb{P}}} \exp((\lambda_- - \lambda_-^{\mathbb{P}})x) & x < 0 \quad y = 0 \\ \frac{c_t^-}{c_t^{\mathbb{P}-}} \frac{\lambda_-}{\lambda_-^{\mathbb{P}}} \exp((\lambda_- - \lambda_-^{\mathbb{P}})y) & x = 0 \quad y < 0 \end{cases}.$$

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 whether we need to specify a maximum flexible pricing kernel where all the parameters for jump intensities are shifted from the \mathbb{Q} - to the \mathbb{P} -measure by different amounts, or whether there is a parsimonious specification (one that imposes more restrictions between the \mathbb{Q} - and \mathbb{P} -dynamics) that is able to simultaneously replicate the observed dynamics 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, for example, Eraker (2004), Santa-Clara and Yan (2010), Christoffersen et al. (2012), and Bates (2012).
3. RND2: Negative jumps are not priced $\iff \lambda_-^{\mathbb{P}} = \lambda_-$, $c_j^{\mathbb{P}-} = c_j^-$, for $j = 0, 1, 2, 3$.
4. RND3: Positive jumps are not priced $\iff \lambda_+^{\mathbb{P}} = \lambda_+$, $c_j^{\mathbb{P}+} = c_j^+$, for $j = 0, 1, 2, 3$.

4 Estimation

We largely rely on the recent paper by 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 dynamics 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 equations, which relate the risk-neutral variances and third-order cumulants to the factors; and (2) non-linear equations, 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 $CUM_t^{(n)\mathbb{Q}} = (CUM_{t,\tau_1}^{(n)\mathbb{Q}}, \dots, CUM_{t,\tau_J}^{(n)\mathbb{Q}})^\top$, where $n \in \{2, 3\}$. We further stack the second and third cumulant vector in $CUM_t^{\mathbb{Q}} = (CUM_t^{(2)\mathbb{Q}}, CUM_t^{(3)\mathbb{Q}})^\top$ to build a $2J \times 1$ vector. This implies the following linear measurement equation:

$$CUM_t^{\mathbb{Q}} = \Gamma_0^{cum} + \Gamma_1^{cum} V_t + \Omega_{cum}^{1/2} \vartheta_t^{cum}, \quad (20)$$

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 Eq. (20) 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 equation for the three factors in the AFT model is

$$V_{t+1} = \Phi_0 + \Phi_1 V_t + \Sigma(V_t)^{1/2} \varepsilon_{t+1}, \quad (21)$$

where Φ_0 , Φ_1 , and $\Sigma(V_t)$ are functions of model parameters. They are given in the Internet Appendix (Section A12) to save space. The system (20)-(21) 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 $V_{0|0} = -(K_1^{\mathbb{P}})^{-1} K_0^{\mathbb{P}}$ and $vec(P_{0|0}) = (I_9 - \Phi_1 \otimes \Phi_1)^{-1} vec(\Sigma(V_{0|0}))$, 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:

$$\begin{cases} V_{t+1|t} = \Phi_0 + \Phi_1 V_{t|t} \\ P_{t+1|t} = \Phi_1 P_{t|t} \Phi_1^\top + \Sigma(V_{t|t}) \\ CUM_{t+1|t}^{\mathbb{Q}} = \Gamma_0 + \Gamma_1 V_{t+1|t} \\ M_{t+1|t} = \Gamma_1 P_{t+1|t} \Gamma_1^\top + \Omega_{cum} \end{cases} \quad (22)$$

and the updating step:

$$\begin{cases} V_{t+1|t+1} = \left[V_{t+1|t} + P_{t+1|t} \Gamma_1^\top M_{t+1|t}^{-1} \left(CUM_{t+1}^{\mathbb{Q}} - CUM_{t+1|t}^{\mathbb{Q}} \right) \right]_+ \\ P_{t+1|t+1} = P_{t+1|t} - P_{t+1|t} \Gamma_1^\top M_{t+1|t}^{-1} \Gamma_1 P_{t+1|t}, \end{cases} \quad (23)$$

where $[V]_+$ returns a vector whose i^{th} element is $\max(V_i, 0)$. 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:

$$Lik^{CUM} = -\frac{1}{2} \sum_{t=1}^T \left[\ln \left((2\pi)^{2J} \det(M_{t|t-1}) \right) + \xi_{t,cum}^\top M_{t|t-1}^{-1} \xi_{t,cum} \right], \quad (24)$$

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 Eq. (8) to compute the model-implied $\mu_2^{\mathbb{Q}+}(t, \tau)$ and $\mu_2^{\mathbb{Q}-}(t, \tau)$. Note that $\mathbb{E}_t^{\mathbb{Q}}[r_{t,t+\tau}^2] = CUM_{t,\tau}^{(2)\mathbb{Q}} + \left(CUM_{t,\tau}^{(1)\mathbb{Q}} \right)^2$ 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) to approximate $\Lambda^{\mathbb{Q}}(t, \tau)$ analytically. All the details are provided in Section A7 of the Internet Appendix. Hence, both $\mu_2^{\mathbb{Q}+}(t, \tau)$ and $\mu_2^{\mathbb{Q}-}(t, \tau)$ are non-linear functions of the factor V_t . We construct a Gaussian quasi log-likelihood for the truncated moments $Tmom_t$,

$$Lik^{Tmom} = -\frac{1}{2} \sum_{t=1}^T \left[\ln \left((2\pi)^{2J} \det(\Omega_{Tmom}) \right) + \xi_{t,Tmom}^\top \Omega_{Tmom}^{-1} \xi_{t,Tmom} \right], \quad (25)$$

where $\xi_{t,Tmom} = Tmom_t^{(Obs)} - Tmom_t(V_{t|t})$, Ω_{Tmom} denotes the measurement error variance, $V_{t|t}$ is obtained through the filtering procedure (see Eq. 22 and 23),

$$Tmom_t^+ = \left(\mu_{2,\tau_1}^+, \dots, \mu_{2,\tau_J}^+ \right)^\top; Tmom_t^- = \left(\mu_{2,\tau_1}^-, \dots, \mu_{2,\tau_J}^- \right)^\top; Tmom_t \equiv \left(Tmom_t^{+\top}, Tmom_t^{-\top} \right)^\top,$$

and $Tmom_t^{(Obs)}$ is the time- t observed risk-neutral truncated moments (computed model-free using Eq. 13). Parameters for different models are estimated via a maximization of $Lik^{CUM} + Lik^{Tmom}$.

4.3 Discussions

Given the heteroscedasticity and non-normality of factors and some nonlinear measurement equations, the Kalman filter is not optimal in this case.⁶ Two other alternatives could have been considered regarding nonlinear measurement equations: (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). We want to emphasize that the Kalman filter described in this paper uses nonlinear equations for parameter estimation (see Eq. 25), which implies that nonlinear equations affect the filtering of unobserved factors in an indirect way. While the extended Kalman filter is appealing as it allows the nonlinear measurement equations to directly affect the filtered factors, in our case its implementation is difficult mainly because nonlinear measurements and their derivatives are approximate.

Nevertheless, of first importance is evaluating the impact of this estimation choice on our main results. To do so, we implement the extended Kalman filter for the most flexible model (AFT4) and compare its fit with the classic Kalman filter. Using the approach in Fang and Oosterlee (2008), we derive a simple analytical approximation of the Jacobian required for the effective implementation of the extended Kalman filter. Because of space constraints, we report these derivations and all empirical results in the Internet Appendix (Table A3, Figures A5 and A6). We find that the two methods provide a very similar fit. Thus, our main conclusions are robust to the choice of the estimation method.

⁶Regarding the issue of heteroscedastic and non-Gaussian factors, we refer readers to Duan and Simonato (1999) for extensive discussions and Monte Carlo analyses suggesting that the loss of optimality is very minimal.

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 a symmetric jump distribution and AFT4 in Andersen et al. (2015). Finally, we evaluate the ability of different pricing kernels with various flexibility to fit 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):

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

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 3.⁷ Overall, regarding the fitting of the term structure of the risk-neutral expected gain and loss, we find that the benchmark two-factor 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 the horizon and ranges from 1% at 2 months to 2.15% at 1 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 a 56% improvement over the benchmark AFT0 model. This performance of the AFT4 model is robust across horizons, with an RMSE as low as 0.45% around horizons of 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 variant's RMSEs. The best-performing model on this front is the AFT1 model with an average RMSE of 0.98%, while the performance of the AFT2, AFT3, and

⁷All risk-neutral parameter estimates, together with their standard deviations, are reported in Table A2 of the Internet Appendix.

AFT4 models is similar. However, the AFT0 model fits the term structure of the total risk-neutral 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 an 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 the risk-neutral variance.

On the term structure of the risk-neutral skewness dimension, the benchmark AFT0 is the worst performer with an average RMSE of 0.85, whereas all the other variants have a similar fit, with an average RMSE of approximately 0.15, which is an almost 80% improvement over the benchmark AFT0. These results underscore the importance of jumps when fitting the term structure of risk-neutral skewness. To better understand our findings, in Figure 2 we plot the observed and model implied average term structure of risk-neutral moments (the top and middle panels). 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; accounting for only the leverage effect is not enough.

The ranking between two-factor models (AFT1 and AFT2) is mixed. The model without a pure-jump process (AFT1) dominates the one without a 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 2 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 the 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-factor models (AFT3 and AFT4), we evaluate the importance of introducing a wedge between the negative and positive jump distributions. 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 has a better fit on the the short end of the risk-neutral expected quadratic gain, while the AFT3 is preferred on the long end.

5.2 Fitting the Quadratic Risk Premium

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. In the bottom panels of Figure 2, we plot the observed and model 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 that is able to adequately fit 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 in 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 term structures, it falls short in capturing the level. Overall, we find that it is imperative to price jumps asymmetrically in the pricing kernel.

To confirm these visual findings, we report the RMSE in Table 4. In addition to the loss and gain QRP RMSEs reported in the top panels, we also report the total QRP (denoted Net QRP since it is the difference between the loss and gain QRP) in the bottom left panel and the skewness risk premium (SRP, 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 visual findings of Figure 2. Except for the very short maturity (three months), $RND4$ yields the smallest RMSE across maturities and for different types of risk premia. $RND0$ is the worst performer, which implies that pricing jumps is important for the dynamics 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 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 the one hand, the *RND4* model offers approximately 70%, 30%, 80%, and 72% improvements over *RND0* for the fitting of the loss, gain, and net QRP, and the SRP, 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, and net QRP, and the SRP, 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 QRP, and the SRP, 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 that 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 the loss and gain quadratic risk premium.

References

- Ait-Sahalia, Yacine, Karaman, Mustafa, and Mancini, Lorian (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, 1926–2010,” *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 Ornathanalai, Chayawat (2012), “Dynamic jump intensities and risk premiums: Evidence from S&P 500 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, I., Giglio, S., Le, A., and Rodriguez, M. (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, B., Jahan-Parvar, M. R., and Tédongap, R. (2013), “Modeling Market Downside Volatility,” *Review of Finance*, 17(1), 443–481.
- 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.
- 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 ($\mathbb{E}^Q[r^2]$, $\mathbb{E}^Q[l^2]$, $\mathbb{E}^Q[g^2]$), the expected quadratic payoff ($\mathbb{E}[r^2]$, $\mathbb{E}[l^2]$, $\mathbb{E}[g^2]$), 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 reported statistics are monthly squared percentage values. The sample period is from January 1996 to December 2015.

Maturity	Mean											
	1	2	3	4	5	6	7	8	9	10	11	12
$\mathbb{E}[r^2]$	26.28 (2.49)	26.49 (2.07)	27.42 (2.07)	28.20 (2.01)	27.97 (1.66)	27.88 (1.45)	27.97 (1.34)	28.06 (1.25)	28.06 (1.16)	28.16 (1.09)	28.38 (1.04)	28.64 (1.00)
$\mathbb{E}[l^2]$	10.87 (1.58)	10.05 (1.30)	10.05 (1.54)	9.95 (1.54)	9.09 (1.13)	8.35 (0.83)	7.90 (0.73)	7.43 (0.66)	6.95 (0.55)	6.56 (0.47)	6.32 (0.44)	6.14 (0.42)
$\mathbb{E}[g^2]$	15.41 (1.07)	16.45 (0.95)	17.37 (0.87)	18.25 (0.88)	18.88 (0.89)	19.53 (0.91)	20.07 (0.89)	20.63 (0.86)	21.11 (0.84)	21.59 (0.82)	22.06 (0.80)	22.50 (0.78)
$\mathbb{E}^Q[r^2]$	45.30 (3.53)	45.55 (3.03)	46.47 (2.92)	47.07 (2.66)	47.39 (2.43)	47.86 (2.38)	48.40 (2.34)	48.81 (2.32)	49.14 (2.29)	49.33 (2.26)	49.40 (2.19)	49.94 (2.18)
$\mathbb{E}^Q[l^2]$	30.46 (2.58)	31.16 (2.26)	32.01 (2.24)	32.50 (2.04)	32.71 (1.85)	33.03 (1.84)	33.39 (1.84)	33.65 (1.84)	33.81 (1.85)	33.83 (1.83)	33.74 (1.77)	34.09 (1.78)
$\mathbb{E}^Q[g^2]$	14.84 (0.98)	14.39 (0.80)	14.46 (0.73)	14.57 (0.68)	14.68 (0.64)	14.84 (0.61)	15.01 (0.59)	15.16 (0.58)	15.33 (0.57)	15.50 (0.56)	15.66 (0.56)	15.85 (0.55)
QRP	19.04 (1.35)	19.07 (1.43)	19.07 (1.43)	18.89 (1.36)	19.43 (1.41)	20.01 (1.44)	20.44 (1.41)	20.77 (1.43)	21.10 (1.46)	21.19 (1.45)	21.04 (1.43)	21.32 (1.44)
QRP^l	19.61 (1.42)	21.13 (1.54)	21.98 (1.63)	22.57 (1.65)	23.64 (1.69)	24.70 (1.71)	25.50 (1.69)	26.23 (1.71)	26.89 (1.74)	27.29 (1.73)	27.43 (1.69)	27.97 (1.70)
QRP^g	0.56 (0.33)	2.05 (0.37)	2.91 (0.40)	3.68 (0.46)	4.20 (0.44)	4.69 (0.46)	5.06 (0.47)	5.46 (0.47)	5.78 (0.47)	6.10 (0.47)	6.40 (0.47)	6.65 (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 (%)			
$\mathbb{E}[l^2]$ and $\mathbb{E}[g^2]$	58.01	37.10	3.13	98.24
$\mathbb{E}^Q[l^2]$ and $\mathbb{E}^Q[g^2]$	87.33	8.78	2.76	98.87
QRP^l and QRP^g	73.05	15.22	6.18	94.45
$\mathbb{E}[l^2]$, $\mathbb{E}[g^2]$, 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

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

where Mom_t^{Mkt} is the time t risk-neutral moment value observed on the market, and 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

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

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. Net QRP is the difference between loss and gain QRP, while skewness RP is the sum. 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					Skewness RP				
	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 the top two panels we plot the average S&P 500 expected quadratic payoff and its loss and gain components for maturities of 1, 3, 6, 9, and 12 months. The top left is for the risk-neutral quantities, while the top right is for the physical quantities. In the middle panels, we plot the level (6-month maturity) and the slope (12-month minus the 2-month maturity) of the risk-neutral expected quadratic payoff and its loss and gain components. In the bottom panels, we plot the physical level and slope. All reported values are monthly and in squared percentage units. The sample period is from January 1996 to December 2015.

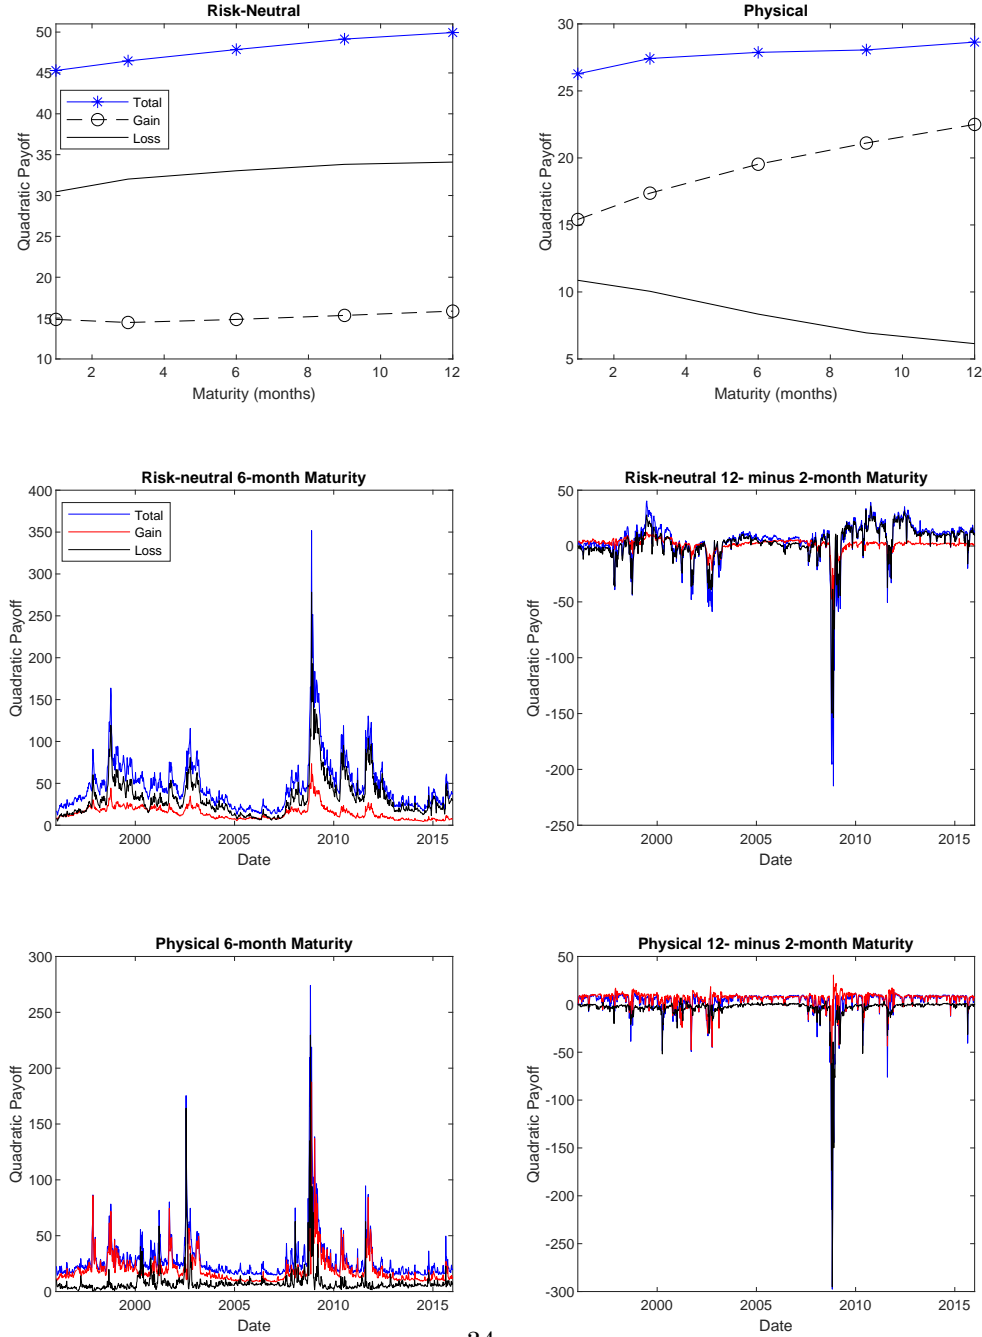



Figure 2: Observed and model-implied average term structure of risk-neutral moments and risk-premium

In this figure we plot the observed and model-implied average term structure of risk-neutral moments (the first two rows) and quadratic risk premium (the third row). All values of the risk-neutral expected quadratic payoff and its components are reported in annualized volatility terms. 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.



Internet Appendix to Accompany “The Term Structures of
Expected Loss and Gain Uncertainty”

March 2020

This appendix contains additional results that are omitted from the main text for brevity.

Contents

A Additional Results	1
A.1 Quadratic Risk Premium Relation with the Variance Risk Premium	1
A.2 Term Structure of Forward Variance Prices	3
A.3 Term Structure of Quadratic Risk Premium	4
A.4 $\eta = 0$ vs $\eta \neq 0$	5
A.5 Kalman Filtering vs Extended Kalman Filtering	5
A.6 Risk-Neutral Moments of Gain and Loss from OTM Options	6
A.7 Computing the Expected Quadratic Loss and Gain using Fang and Oosterlee (2008)	10
A.8 Computing the Jacobian of the Expected Quadratic Loss and Gain using Fang and Oosterlee (2008)	12
A.9 Physical Expected Quadratic Payoff and the Drift	15
A.10 Normality Assumption and Physical Expected Quadratic Loss and Gain	15
A.11 Rolling Window Parameter Estimation	18
A.12 Transition Equation for the Three Factors in the Andersen et al. (2015b) Model . . .	20

A Additional Results

A.1 Quadratic Risk Premium Relation with the Variance Risk Premium

In this section, we follow Feunou et al. (2019) and discuss in detail the difference between the quadratic risk premium (QRP) and the variance risk premium (VRP). Both QRP and VRP share the premium definition, but they regard different measures of uncertainty: the quadratic payoff versus the realized variance. Therefore, the difference between QRP and VRP hinges on the difference between the quadratic payoff and the realized variance. For a given stock, we observe returns at regular high-frequency time intervals of length δ . The monthly realized return $r_{t-1,t}$ and the monthly realized variance $RV_{t-1,t}$ are defined by aggregating $r_{t-1+j\delta}$ and $r_{t-1+j\delta}^2$, respectively:

$$r_{t-1,t} = \sum_{j=1}^{1/\delta} r_{t-1+j\delta} \quad \text{and} \quad RV_{t-1,t} = \sum_{j=1}^{1/\delta} r_{t-1+j\delta}^2, \quad (\text{A.1})$$

where $1/\delta$ is the number of high-frequency returns in a monthly period, e.g., $\delta = 1/21$ for daily returns, and $r_{t-1+j/21}$ denotes the j th high-frequency return of the monthly period starting from day $t-1$ and ending on day t . The quadratic payoff and the realized variance are related as follows:

$$r_{t-1,t}^2 = RV_{t-1,t} + 2RA_{t-1,t}, \quad \text{where} \quad RA_{t-1,t} = \sum_{i=1}^{1/\delta-1} \sum_{j=1}^{1/\delta-i} r_{t-1+j\delta} r_{t-1+j\delta+i\delta}, \quad (\text{A.2})$$

and $RA_{t-1,t}$ is the realized autocovariance.

The realized variance is a measure of fluctuating uncertainty based on higher-frequency returns, while the quadratic payoff is a measure of fluctuating uncertainty based on lower-frequency returns. Equation (A.2) shows that the quadratic payoff is approximately equal to the realized variance if and only if the realized autocovariance is negligible. Feunou et al. (2019) show that using either daily or intra-day data, the monthly realized autocovariance is not negligible.

To study the difference between QRP and VRP, we adopt the theoretical definition of VRP in

Bollerslev et al. (2009) as follows:

$$\begin{aligned} VRP_t &\equiv \mathbb{E}_t^{\mathbb{Q}}[RV_{t,t+1}] - \mathbb{E}_t[RV_{t,t+1}] \\ &= \sum_{j=1}^{1/\delta} \left(\mathbb{E}_t^{\mathbb{Q}}[r_{t-1+j\delta}^2] - \mathbb{E}_t[r_{t-1+j\delta}^2] \right). \end{aligned} \quad (\text{A.3})$$

In the empirical exercises of the VRP literature, the risk-neutral expectation of quadratic payoff $\mathbb{E}_t^{\mathbb{Q}}[r_{t,t+1}^2]$ is often used to proxy for the risk-neutral expected realized variance $\mathbb{E}_t^{\mathbb{Q}}[RV_{t,t+1}]$. This is, for example, the case in Feunou et al. (2017) and Kilic and Shaliastovich (2019). By doing so, they use an empirical measure of the variance risk premium \widetilde{VRP}_t defined by:

$$\begin{aligned} \widetilde{VRP}_t &= \mathbb{E}_t^{\mathbb{Q}}[r_{t,t+1}^2] - \mathbb{E}_t[RV_{t,t+1}] \\ &= VRP_t + 2\mathbb{E}_t^{\mathbb{Q}}[RA_{t,t+1}]. \end{aligned} \quad (\text{A.4})$$

By definition, \widetilde{VRP}_t is not a coherent measure of a risk premium (i.e., it is not the difference between the risk-neutral and physical expectations of the same quantity). Instead, \widetilde{VRP}_t is a biased measure of VRP_t , where the bias equals $2\mathbb{E}_t^{\mathbb{Q}}[RA_{t,t+1}]$. Feunou et al. (2019) show that this bias is not negligible for the S&P 500.

Lastly, the bias from the realized autocovariance is even more severe when we decompose VRP into its loss and gain components. The gain-loss decomposition of the squared return allows us to write the realized variance as the total of two components: the cumulative sum of squared high-frequency gains and the cumulative sum of squared high-frequency losses, which, similar to the quadratic gain and the quadratic loss can be interpreted as measures of gain uncertainty and loss uncertainty, respectively. These two components of the realized variance are what Barndorff-Nielsen et al. (2010) refer to as realized semi-variances, formally defined as:

$$RV_{t-1,t} = RV_{t-1,t}^g + RV_{t-1,t}^l \quad \text{where} \quad RV_{t-1,t}^g = \sum_{j=1}^{1/\delta} g_{t-1+j\delta}^2 \quad \text{and} \quad RV_{t-1,t}^l = \sum_{j=1}^{1/\delta} l_{t-1+j\delta}^2, \quad (\text{A.5})$$

where $RV_{t-1,t}^g$ and $RV_{t-1,t}^l$ are referred to as bad and good variances in the literature (e.g., Patton and Sheppard 2015; Kilic and Shaliastovich 2019). Note that, even if a negligible magnitude of the realized autocovariance made the quadratic payoff $r_{t-1,t}^2$ proxy for the realized variance $RV_{t-1,t}$, the

quadratic loss (gain) would not proxy for the loss (gain) realized variance. In fact, the quadratic loss (gain) is a censored variable while the loss (gain) realized variance is not. Feunou et al. (2019) show that the S&P 500 quadratic loss (gain) and the loss (gain) realized variance have very different dynamics.

In summary, we show that the quadratic payoff can be very different from the realized variance. This is also partly because the equity risk premium at a lower frequency is non-zero and time-varying. The empirical evidence in Feunou et al. (2019) supports this large wedge between the quadratic payoff and the realized variance at a monthly frequency. The difference is much more significant between the quadratic loss and gain and their corresponding semi-variances.

A.2 Term Structure of Forward Variance Prices

Dew-Becker et al. (2017) compute the term structure of forward variance prices as $F_t^{rv,\tau} \equiv \mathbb{E}_t^{\mathbb{Q}} [RV_{t+\tau-1,t+\tau}]$. 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 $t + \tau$. Dew-Becker et al. compute these forward prices using traded variance swaps. Since traded loss and gain variance swaps do not exist, one cannot estimate risk-neutral expectations of the loss and gain components using their approach. Previous literature (see, e.g., Kilic and Shaliastovich 2019) has shown that loss and gain components are important for asset prices. Our methodology allows us to compute the term structure of forward prices for not only the quadratic payoff but also its loss and gain components as $F_t^{r,\tau} \equiv \mathbb{E}_t^{\mathbb{Q}} [r_{t+\tau-1,t+\tau}^2]$, $F_t^{l,\tau} \equiv \mathbb{E}_t^{\mathbb{Q}} [l_{t+\tau-1,t+\tau}^2]$ and $F_t^{g,\tau} \equiv \mathbb{E}_t^{\mathbb{Q}} [g_{t+\tau-1,t+\tau}^2]$, respectively. The risk-neutral expectations of the loss and gain components are computed following Equations (6) and (7) in the main paper.

In Panel A of Figure A1, we plot the term structure of average forward prices for the risk-neutral expected quadratic payoff and its loss and gain components. The average forward prices are reported in annualized percentage volatility units: $100 \times \sqrt{12 \times F_t^{\tau}}$. We find that the estimated term structure of forward prices for the quadratic payoff is concave with both the level and slope similar to the forward prices in Dew-Becker et al. This is despite the differences in our sample periods, and the fact that our forward prices are computed for the quadratic payoff while their forward prices are computed from traded variance swaps. Most importantly, as mentioned above, in

contrast to Dew-Becker et al. we are able to separately estimate the forward prices for the quadratic loss and gain while loss and gain variance swaps do not exist. This allows us to investigate not only the term structure of total quadratic payoff but also its components. We see that the term structure of average forward prices for the quadratic loss and gain are in general upward sloping. We also find that, across all horizons, the quadratic loss forward prices are higher than the quadratic gain forward prices.

A.3 Term Structure of Quadratic Risk Premium

In Panel B of Figure A1, we plot the time-series average of the term structure of quadratic risk premium (QRP) and its loss and gain components for maturities of 1, 3, 6, 9 and 12 months. The loss QRP is the price investors pay to hedge extreme losses during bad times. On the other hand, the gain QRP is the compensation investors demand for weak upside potential during bad times. We find that the loss, gain and net QRP are all positive across all maturities. This suggests that regardless of the investment horizon, investors are always averse to fluctuating loss or gain uncertainty. Furthermore, we find that the term structures of the loss and gain QRP are upward sloping. This suggests that investors are willing to pay (demand to earn) a higher premium to insure against strong downside risk (weak upside potential) in the long run. The term structure of net QRP is also upward sloping. Since the net QRP represents the net cost of insuring fluctuations in loss uncertainty, the observed term structure pattern suggests that this cost is increasing with the investment horizon.

In Figure A2, we plot the 6-month (the level) and the 12-month minus 2-month (the slope) maturity of the QRP (net, loss and gain). In Panel A, we see that the level of the loss QRP is positive in 98% of the sample. This implies that regardless of the economic conditions, investors always seek to hedge extreme losses. The level of loss QRP spikes in periods of high uncertainty. The slopes of the loss QRP fluctuate and are, in general, positive (75% of the sample), but become negative, especially during crises. This implies that during calm periods, investors are willing to pay more to hedge against downside risk in the long run versus the short run. However, during crises they are much more concerned about hedging short-term downside risks. From Panel B, we also see that the level and the slope of the gain QRP are positive in general (85% and 88% of the sample, respectively). Hence, investors are averse to fluctuating quadratic gain and demand a

higher premium for fluctuations in gain uncertainty in the long run versus the short run.

A.4 $\eta = 0$ vs $\eta \neq 0$

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.¹ We report factors estimates in Figure A3, the model-implied average risk-neutral moments in Figure A4 and risk-neutral moments RMSEs in Table A1.

From a pure statistical standpoint, the restriction $\eta = 0$ is rejected at any significance level (indeed, the average likelihood for the unrestricted case is 6.25, while imposing $\eta = 0$ substantially reduces it to 5.47). As shown in Figure A3, the filtered third factor $V_{3,t}$ is a lot more volatile when imposing $\eta = 0$. It is not surprising that imposing η makes $V_{3,t}$ a pure jump process in the sense that it has no impact in the diffusive volatility of returns. On average, as shown in Figure A4, the restriction $\eta = 0$ is unable to capture the term structure of risk-neutral moments. This is confirmed in Table A1 where we report the RMSEs across models, moments and maturities. Except for risk-neutral skewness, the restriction $\eta = 0$ produces RMSE numbers that are significantly higher than the unrestricted case.

A.5 Kalman Filtering vs Extended Kalman Filtering

The estimation procedure described in the main text uses only linear measurement equations to filter the latent state V_t . In this section we evaluate the robustness of our findings to the estimation procedure. To do that, we implement the extended Kalman filter, which differs from the classic Kalman by using all the measurement equations (linear and nonlinear) in the filtering step. The linear measurement equations relate risk-neutral cumulant to V_t ,

$$CUM_t^{\mathbb{Q}} = \Gamma_0 + \Gamma_1 V_t,$$

while nonlinear measurement equations relate risk-neutral expected quadratic negative and positive returns to V_t ,

$$Tmom_t^{\mathbb{Q}} = h(V_t).$$

¹Several contributions including Santa-Clara and Yan (2010), Christoffersen et al. (2012), and Andersen et al. (2015b) find evidence for a pure jump component in the pricing of S&P 500 options.

Let us denote by $H(\cdot)$ the Jacobian of the function $h(\cdot)$; the extended Kalman filter proceeds recursively through the forecasting step:

$$\left\{ \begin{array}{l} V_{t+1|t} = \Phi_0 + \Phi_1 V_{t|t} \\ P_{t+1|t} = \Phi_1 P_{t|t} \Phi_1^\top + \Delta t \Sigma(V_{t|t}) \\ CUM_{t+1|t}^{\mathbb{Q}} = \Gamma_0 + \Gamma_1 V_{t+1|t} \\ Tmom_{t+1|t}^{\mathbb{Q}} = h(V_{t+1|t}) \\ \Gamma_{t+1|t} = \left[\Gamma_1^\top, H(V_{t+1|t})^\top \right]^\top \\ M_{t+1|t} = \Gamma_{t+1|t} P_{t+1|t} \Gamma_{t+1|t}^\top + \Omega_e, \end{array} \right. \quad (\text{A.6})$$

and the updating step:

$$\left\{ \begin{array}{l} V_{t+1|t+1} = \left[V_{t+1|t} + P_{t+1|t} \Gamma_{t+1|t}^\top M_{t+1|t}^{-1} \left(MOM_{t+1}^{\mathbb{Q}} - MOM_{t+1|t}^{\mathbb{Q}} \right) \right]_+, \\ P_{t+1|t+1} = P_{t+1|t} - P_{t+1|t} \Gamma_{t+1|t}^\top M_{t+1|t}^{-1} \Gamma_{t+1|t} P_{t+1|t}, \end{array} \right. \quad (\text{A.7})$$

where $MOM_t^{\mathbb{Q}} \equiv \left(CUM_t^{\mathbb{Q}\top}, Tmom_t^{\mathbb{Q}\top} \right)^\top$. We provide all the details for the computation of function $H(\cdot)$ in Section A.8 of this appendix.

The comparison between the two estimation methods is done for the most flexible model (AFT4). We report factors estimates in Figure A5, implied average risk-neutral moments in Figure A6 and risk-neutral moments RMSEs in Table A3.

The extended Kalman filter generates an average likelihood of 6.27, which is slightly higher than that of the classic Kalman filter (6.25); the difference is statistically meaningless. To confirm this similarity, we plot the filtered factors in Figure A5, where it is obvious that the two set of factors are almost indistinguishable. On average, as shown in Figure A6, the two estimation techniques generate similar average term structures of risk-neutral moments. This is confirmed in Table A3, where we report the RMSEs across estimation methods, moments and maturities. It is readily apparent that the two methods generate similar fit for all the risk-neutral moments.

A.6 Risk-Neutral Moments of Gain and Loss from OTM Options

In this section, we present the proof established in Feunou et al. (2019) to show that $V_t^g(\tau)$ is the price of the quadratic gain, and therefore $V_t^l(\tau)$ is the price of the quadratic loss. Consider the

function

$$F(X) = \frac{1}{\alpha} \ln(1 - \delta + \delta \exp(\alpha X))$$

with $0 \leq \delta \leq 1$ and $\alpha > 0$. It can easily be verified that $F(X) = \max(X, 0)$ if $\alpha \rightarrow \infty$, $0 < \delta < 1$.

Suppose we are interested in computing the risk-neutral moments of the gain component of the τ -period log returns defined by $r_{t,t+\tau} = \ln \left[\frac{S_{t+\tau}}{S_t} \right]$. That is, we want to compute

$$\mathbb{E}_t^{\mathbb{Q}} [g_{t,t+\tau}^n] \quad \text{for } n \geq 2 \quad \text{where } g_{t,t+\tau} = \max(r_{t,t+\tau}, 0).$$

Observe that

$$g_{t,t+\tau}^n = (\max(r_{t,t+\tau}, 0))^n = \lim_{\substack{\alpha \rightarrow \infty \\ 0 < \delta < 1}} (F(r_{t,t+\tau}))^n.$$

It follows that

$$\mathbb{E}_t^{\mathbb{Q}} [g_{t,t+\tau}^n] = \lim_{\substack{\alpha \rightarrow \infty \\ 0 < \delta < 1}} \mathbb{E}_t^{\mathbb{Q}} [(F(r_{t,t+\tau}))^n] \quad \text{for } n \geq 2. \quad (\text{A.8})$$

Remark that $F(0) = 0$ and that F is twice differentiable with

$$\begin{aligned} F'(X) &= \frac{\delta \exp(\alpha X)}{1 - \delta + \delta \exp(\alpha X)} = \delta \exp(\alpha(X - F(X))) \\ F''(X) &= \delta \alpha (1 - F'(X)) \exp(\alpha(X - F(X))) = \alpha (1 - F'(X)) F'(X) = \frac{\alpha \delta (1 - \delta) \exp(\alpha X)}{(1 - \delta + \delta \exp(\alpha X))^2}. \end{aligned}$$

Thus we can compute $\mathbb{E}_t^{\mathbb{Q}} [(F(r_{t,t+\tau}))^n]$ for $n \geq 2$ by applying the Bakshi et al. (2003) formula

$$\begin{aligned} \mathbb{E}_t^{\mathbb{Q}} [\exp(-r\tau) H(S_{t+\tau})] &= \exp(-r\tau) (H(S_t) - S_t H'(S_t)) + S_t H'(S_t) \\ &\quad + \int_0^{S_t} H''(K) P(t, \tau; K) dK + \int_{S_t}^{\infty} H''(K) C(t, \tau; K) dK \end{aligned} \quad (\text{A.9})$$

with the twice differentiable function $H(S) = \left(F \left(\ln \left[\frac{S}{S_t} \right] \right) \right)^n$.

We have

$$H'(S) = \frac{n F' \left(\ln \left[\frac{S}{S_t} \right] \right) \left(F \left(\ln \left[\frac{S}{S_t} \right] \right) \right)^{n-1}}{S}$$

and

$$H''(S) = \frac{n \left[\left(F'' \left(\ln \left[\frac{S}{S_t} \right] \right) - F' \left(\ln \left[\frac{S}{S_t} \right] \right) \right) F \left(\ln \left[\frac{S}{S_t} \right] \right) + (n-1) \left(F' \left(\ln \left[\frac{S}{S_t} \right] \right) \right)^2 \right] \left(F \left(\ln \left[\frac{S}{S_t} \right] \right) \right)^{n-2}}{S^2}.$$

Observe that, since $F(0) = 0$ and $F'(0) = \delta$, for $n \geq 2$ we have

$$H(S_t) = (F(0))^n = 0 \quad \text{and} \quad H'(S_t) = \frac{nF'(0)(F(0))^{n-1}}{S_t} = 0.$$

This means that

$$\exp(-r\tau) (H(S_t) - S_t H'(S_t)) + S_t H'(S_t) = 0. \quad (\text{A.10})$$

Now, we are interested in computing

$$\lim_{\substack{\alpha \rightarrow \infty \\ 0 < \delta < 1}} H''(K).$$

We have

$$H''(K) = \frac{n \left[(F''(X) - F'(X)) F(X) + (n-1) (F'(X))^2 \right] (F(X))^{n-2}}{K^2} \quad \text{where} \quad X = \ln \left[\frac{K}{S_t} \right].$$

For OTM put options, we have $K < S_t$ or equivalently $X < 0$. Observe from their expressions that when $\alpha \rightarrow \infty$, $0 < \delta < 1$, then $F(X) \rightarrow \max(X, 0) = 0$, $F'(X) \rightarrow 0$ and also $F''(X) \rightarrow 0$.

This means that

$$\forall K < S_t \quad \lim_{\substack{\alpha \rightarrow \infty \\ 0 < \delta < 1}} H''(K) = 0$$

and thus

$$\begin{aligned} \lim_{\substack{\alpha \rightarrow \infty \\ 0 < \delta < 1}} \int_0^{S_t} H''(K) P(t, \tau; K) dK &= \int_0^{S_t} \left(\lim_{\substack{\alpha \rightarrow \infty \\ 0 < \delta < 1}} H''(K) \right) P(t, \tau; K) dK \\ &= 0. \end{aligned} \quad (\text{A.11})$$

For OTM call options, we have $K > S_t$ or equivalently $X > 0$. Observe from their expressions that when $\alpha \rightarrow \infty$, $0 < \delta < 1$, then $F(X) \rightarrow \max(X, 0) = X$, $F'(X) \rightarrow 1$ and $F''(X) \rightarrow 0$. This

means that

$$\forall K > S_t \quad \lim_{\substack{\alpha \rightarrow \infty \\ 0 < \delta < 1}} H''(K) = \frac{n \left(n - 1 - \ln \left[\frac{K}{S_t} \right] \right) \left(\ln \left[\frac{K}{S_t} \right] \right)^{n-2}}{K^2}$$

and thus

$$\begin{aligned} \lim_{\substack{\alpha \rightarrow \infty \\ 0 < \delta < 1}} \int_{S_t}^{\infty} H''(K) C(t, \tau; K) dK &= \int_{S_t}^{\infty} \left(\lim_{\substack{\alpha \rightarrow \infty \\ 0 < \delta < 1}} H''(K) \right) C(t, \tau; K) dK \\ &= \int_{S_t}^{\infty} \frac{n \left(n - 1 - \ln \left[\frac{K}{S_t} \right] \right) \left(\ln \left[\frac{K}{S_t} \right] \right)^{n-2}}{K^2} C(t, \tau; K) dK. \end{aligned} \quad (\text{A.12})$$

Taking the limit of Equation (A.9) when $\alpha \rightarrow \infty$, $0 < \delta < 1$, equations (A.10), (A.11) and (A.12) imply that

$$\mathbb{E}_t^{\mathbb{Q}} [\exp(-r\tau) g_{t,t+\tau}^n] = \int_{S_t}^{\infty} \frac{n \left(n - 1 - \ln \left[\frac{K}{S_t} \right] \right) \left(\ln \left[\frac{K}{S_t} \right] \right)^{n-2}}{K^2} C(t, \tau; K) dK \quad \text{for } n \geq 2. \quad (\text{A.13})$$

Since Bakshi et al. (2003) show that

$$\begin{aligned} \mathbb{E}_t^{\mathbb{Q}} [\exp(-r\tau) r_{t,t+\tau}^n] &= \int_0^{S_t} \frac{n \left(n - 1 + \ln \left[\frac{S_t}{K} \right] \right) \left(-\ln \left[\frac{S_t}{K} \right] \right)^{n-2}}{K^2} P(t, \tau; K) dK \\ &\quad + \int_{S_t}^{\infty} \frac{n \left(n - 1 - \ln \left[\frac{K}{S_t} \right] \right) \left(\ln \left[\frac{K}{S_t} \right] \right)^{n-2}}{K^2} C(t, \tau; K) dK \quad \text{for } n \geq 2, \end{aligned} \quad (\text{A.14})$$

and given that $r_{t,t+\tau}^n = g_{t,t+\tau}^n + (-1)^n l_{t,t+\tau}^n$, where $l_{t,t+\tau} = \max(-r_{t,t+\tau}, 0)$, then it follows that

$$\mathbb{E}_t^{\mathbb{Q}} [\exp(-r\tau) l_{t,t+\tau}^n] = \int_0^{S_t} \frac{n \left(n - 1 + \ln \left[\frac{S_t}{K} \right] \right) \left(\ln \left[\frac{S_t}{K} \right] \right)^{n-2}}{K^2} P(t, \tau; K) dK \quad \text{for } n \geq 2. \quad (\text{A.15})$$

A.7 Computing the Expected Quadratic Loss and Gain using Fang and Oosterlee (2008)

We follow Fang and Oosterlee (2008) to approximate $\Lambda^{\mathbb{Q}}(t, \tau)$ analytically as:

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

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

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

and for $k > 0$

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

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 .² Next, we provide all the details leading up to these closed-form expressions.

We can split ω_k into

$$\omega_k = \omega_k^- - \omega_k^+$$

²The 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. 2015b for more details). We thank Nicola Fusari for sharing the estimation code.

where

$$\begin{aligned}\omega_0^- &= \frac{1}{b-a} \int_a^b y^2 \mathbb{I}(y < 0) dy, & \omega_0^+ &= \frac{1}{b-a} \int_a^b y^2 \mathbb{I}(y > 0) dy \\ \omega_k^- &= \frac{2}{b-a} \int_a^b y^2 \mathbb{I}(y < 0) \cos\left(k\pi \frac{y-a}{b-a}\right) dy, & \text{for } k \geq 1 \\ \omega_k^+ &= \frac{2}{b-a} \int_a^b y^2 \mathbb{I}(y > 0) \cos\left(k\pi \frac{y-a}{b-a}\right) dy & \text{for } k \geq 1\end{aligned}$$

We also have

$$\omega_k^- = V_k - \omega_k^+$$

where

$$\begin{aligned}V_0 &= \frac{1}{b-a} \int_a^b y^2 dy, & V_k &= \frac{2}{b-a} \int_a^b y^2 \cos\left(k\pi \frac{y-a}{b-a}\right) dy, & \text{for } k \geq 1 \\ \omega_0^+ &= \frac{1}{b-a} \int_a^b y^2 \mathbb{I}(y > 0) dy = \frac{1}{b-a} \int_{a^+}^b y^2 dy = \frac{1}{3(b-a)} [y^3]_{a^+}^b\end{aligned}$$

and

$$\begin{aligned}\omega_k^+ &= \frac{2}{b-a} \int_{a^+}^b z^2 \cos\left(k\pi \frac{z-a}{b-a}\right) dz \\ &= \frac{2}{b-a} \frac{b-a}{k\pi} \left[z^2 \sin\left(k\pi \frac{z-a}{b-a}\right) \right]_{a^+}^b - \frac{2}{b-a} \frac{b-a}{k\pi} \int_{a^+}^b 2z \sin\left(k\pi \frac{z-a}{b-a}\right) dz \\ &= \frac{2}{k\pi} \left[b^2 \sin\left(k\pi \frac{b-a}{b-a}\right) - (a^+)^2 \sin\left(k\pi \frac{a^+-a}{b-a}\right) \right] - \frac{2}{k\pi} \int_{a^+}^b 2z \sin\left(k\pi \frac{z-a}{b-a}\right) dz \\ &= -\frac{2}{k\pi} (a^+)^2 \sin\left(k\pi \frac{a^+-a}{b-a}\right) + \frac{4}{k\pi} \frac{b-a}{k\pi} \left(\left[z \cos\left(k\pi \frac{z-a}{b-a}\right) \right]_{a^+}^b - \int_{a^+}^b \cos\left(k\pi \frac{z-a}{b-a}\right) dz \right) \\ &= -\frac{2}{k\pi} (a^+)^2 \sin\left(k\pi \frac{a^+-a}{b-a}\right) + \frac{4}{k\pi} \frac{b-a}{k\pi} \left(b \cos\left(k\pi \frac{b-a}{b-a}\right) - a^+ \cos\left(k\pi \frac{a^+-a}{b-a}\right) \right. \\ &\quad \left. - \int_{a^+}^b \cos\left(k\pi \frac{z-a}{b-a}\right) dz \right) \\ &= -\frac{2}{k\pi} (a^+)^2 \sin\left(k\pi \frac{a^+-a}{b-a}\right) + \frac{4}{k\pi} \frac{b-a}{k\pi} \left(b \cos\left(k\pi \frac{b-a}{b-a}\right) - a^+ \cos\left(k\pi \frac{a^+-a}{b-a}\right) \right. \\ &\quad \left. - \frac{b-a}{k\pi} \left[\sin\left(k\pi \frac{z-a}{b-a}\right) \right]_{a^+}^b \right) \\ &= -\frac{2}{k\pi} (a^+)^2 \sin\left(k\pi \frac{a^+-a}{b-a}\right) + \frac{4}{k\pi} \frac{b-a}{k\pi} \left(b(-1)^k - a^+ \cos\left(k\pi \frac{a^+-a}{b-a}\right) + \frac{b-a}{k\pi} \sin\left(k\pi \frac{a^+-a}{b-a}\right) \right) \\ &= \left(\frac{4(b-a)^2}{(k\pi)^3} - \frac{2}{k\pi} (a^+)^2 \right) \sin\left(k\pi \frac{a^+-a}{b-a}\right) - \frac{4a^+(b-a)}{(k\pi)^2} \cos\left(k\pi \frac{a^+-a}{b-a}\right) + \frac{4b(b-a)(-1)^k}{(k\pi)^2}\end{aligned}$$

$$V_0 = \frac{1}{b-a} \int_a^b y^2 dy = \frac{1}{b-a} \int_a^b y^2 dy = \frac{1}{3(b-a)} [y^3]_a^b$$

$$\begin{aligned}
V_k &= \frac{2}{b-a} \int_a^b z^2 \cos \left(k\pi \frac{z-a}{b-a} \right) dz \\
&= \frac{2}{b-a} \frac{b-a}{k\pi} \left[z^2 \sin \left(k\pi \frac{z-a}{b-a} \right) \right]_a^b - \frac{2}{b-a} \frac{b-a}{k\pi} \int_a^b 2z \sin \left(k\pi \frac{z-a}{b-a} \right) dz \\
&= \frac{2}{k\pi} \left[b^2 \sin \left(k\pi \frac{b-a}{b-a} \right) - a^2 \sin \left(k\pi \frac{a-a}{b-a} \right) \right] - \frac{2}{k\pi} \int_a^b 2z \sin \left(k\pi \frac{z-a}{b-a} \right) dz \\
&= -\frac{2}{k\pi} a^2 \sin \left(k\pi \frac{a-a}{b-a} \right) + \frac{4}{k\pi} \frac{b-a}{k\pi} \left(\left[z \cos \left(k\pi \frac{z-a}{b-a} \right) \right]_a^b - \int_a^b \cos \left(k\pi \frac{z-a}{b-a} \right) dz \right) \\
&= -\frac{2}{k\pi} a^2 \sin \left(k\pi \frac{a-a}{b-a} \right) + \frac{4}{k\pi} \frac{b-a}{k\pi} \left(b \cos \left(k\pi \frac{b-a}{b-a} \right) - a \cos \left(k\pi \frac{a-a}{b-a} \right) - \int_{a^+}^b \cos \left(k\pi \frac{z-a}{b-a} \right) dz \right) \\
&= -\frac{2}{k\pi} a^2 \sin \left(k\pi \frac{a-a}{b-a} \right) + \frac{4}{k\pi} \frac{b-a}{k\pi} \left(b \cos \left(k\pi \frac{b-a}{b-a} \right) - a \cos \left(k\pi \frac{a-a}{b-a} \right) - \frac{b-a}{k\pi} \left[\sin \left(k\pi \frac{z-a}{b-a} \right) \right]_a^b \right) \\
&= -\frac{2}{k\pi} a^2 \sin \left(k\pi \frac{a-a}{b-a} \right) + \frac{4}{k\pi} \frac{b-a}{k\pi} \left(b(-1)^k - a \cos \left(k\pi \frac{a-a}{b-a} \right) + \frac{b-a}{k\pi} \sin \left(k\pi \frac{a-a}{b-a} \right) \right) \\
&= \left(\frac{4(b-a)^2}{(k\pi)^3} - \frac{2}{k\pi} a^2 \right) \sin \left(k\pi \frac{a-a}{b-a} \right) - \frac{4a(b-a)}{(k\pi)^2} \cos \left(k\pi \frac{a-a}{b-a} \right) + \frac{4b(b-a)(-1)^k}{(k\pi)^2} \\
&= \frac{4(b-a)b(-1)^k}{(k\pi)^2} - \frac{4(b-a)a}{(k\pi)^2} = \frac{4(b-a)(b(-1)^k - a)}{(k\pi)^2}
\end{aligned}$$

A.8 Computing the Jacobian of the Expected Quadratic Loss and Gain using Fang and Oosterlee (2008)

$$\begin{aligned}
\mu_2^{\mathbb{Q}-}(t, \tau) &= \frac{\mathbb{E}_t^{\mathbb{Q}}[r_{t,t+\tau}^2] + \Lambda^{\mathbb{Q}}(t, \tau)}{2} \\
\mu_2^{\mathbb{Q}+}(t, \tau) &= \frac{\mathbb{E}_t^{\mathbb{Q}}[r_{t,t+\tau}^2] - \Lambda^{\mathbb{Q}}(t, \tau)}{2},
\end{aligned} \tag{A.16}$$

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

where

$$\begin{aligned}
\omega_0 &\equiv \frac{1}{b-a} \int_a^b y^2 \{ \mathbb{I}(y < 0) - \mathbb{I}(y > 0) \} dy \\
\omega_k &\equiv \frac{2}{b-a} \int_a^b y^2 \{ \mathbb{I}(y < 0) - \mathbb{I}(y > 0) \} \cos \left(k\pi \frac{y-a}{b-a} \right) dy, \quad \text{for } k \geq 1.
\end{aligned}$$

Using Fang and Oosterlee (2008), we have established in Section A.7 that

$$\omega_0 = \begin{cases} -\frac{b^3-a^3}{3(b-a)} & \text{if } a \geq 0 \\ -\frac{b^3+a^3}{3(b-a)} & \text{if } a < 0 \end{cases},$$

and for $k > 0$

$$\omega_k = \begin{cases} -\frac{4(b-a)(b(-1)^k-a)}{(k\pi)^2} & \text{if } a \geq 0 \\ -\frac{4(b-a)(b(-1)^k+a)}{(k\pi)^2} + \frac{8(b-a)^2}{(k\pi)^3} \sin \left(k\pi \frac{a}{b-a} \right) & \text{if } a < 0 \end{cases}.$$

Given the structure of the AFT model, we have

$$\varphi_{t,\tau}(u) = \exp \left(\alpha(u, \tau) + \beta(u, \tau)' V_t \right).$$

The first order expansion around a given state \bar{V} is given by

$$\varphi_{t,\tau}(u) = \bar{\varphi}_{t,\tau}(u) + \bar{\varphi}_{t,\tau}(u) \beta(u, \tau)' (V_t - \bar{V}),$$

where $\bar{\varphi}_{t,\tau}(u) = \exp \left(\alpha(u, \tau) + \beta(u, \tau)' \bar{V} \right).$

The derivative of $\Lambda^{\mathbb{Q}}(t, \tau)$ at \bar{V} is given by

$$\frac{\partial \Lambda^{\mathbb{Q}}(t, \tau)}{\partial V_t} \Big|_{V_t = \bar{V}} = \sum_{k=1}^{N-1} \text{Re} \left\{ \bar{\varphi}_{t, \tau} \left(i \frac{k\pi}{b-a} \right) \exp \left(-i \frac{k\pi a}{b-a} \right) \beta \left(i \frac{k\pi}{b-a}, \tau \right)' \right\} \omega_k$$

$$\begin{aligned} \mathbb{E}_t^{\mathbb{Q}} [r_{t, t+\tau}^2] &= \mathbb{V}_t^{\mathbb{Q}} [r_{t, t+\tau}] + \left(\mathbb{E}_t^{\mathbb{Q}} [r_{t, t+\tau}] \right)^2 \\ &= \frac{\partial^2 \alpha(u, \tau)}{\partial u^2} \Big|_{u=0} + \frac{\partial^2 \beta'(u, \tau)}{\partial u^2} \Big|_{u=0} V_t + \left(\frac{\partial \alpha(u, \tau)}{\partial u} \Big|_{u=0} + \frac{\partial \beta'(u, \tau)}{\partial u} \Big|_{u=0} V_t \right)^2 \end{aligned}$$

$$\frac{\partial \mathbb{E}_t^{\mathbb{Q}} [r_{t, t+\tau}^2]}{\partial V_t} \Big|_{V_t = \bar{V}} = \frac{\partial^2 \beta'(u, \tau)}{\partial u^2} \Big|_{u=0} + 2 \left(\frac{\partial \alpha(u, \tau)}{\partial u} \Big|_{u=0} + \frac{\partial \beta'(u, \tau)}{\partial u} \Big|_{u=0} \bar{V} \right) \frac{\partial \beta'(u, \tau)}{\partial u} \Big|_{u=0}$$

$$\begin{aligned} &\frac{\partial \mu_2^{\mathbb{Q}^-}(t, \tau)}{\partial V_t} \Big|_{V_t = \bar{V}} \\ &= \frac{1}{2} \frac{\partial^2 \beta'(u, \tau)}{\partial u^2} \Big|_{u=0} + \left(\frac{\partial \alpha(u, \tau)}{\partial u} \Big|_{u=0} + \frac{\partial \beta'(u, \tau)}{\partial u} \Big|_{u=0} \bar{V} \right) \frac{\partial \beta'(u, \tau)}{\partial u} \Big|_{u=0} \\ &\quad + \frac{1}{2} \left\{ \sum_{k=1}^{N-1} \text{Re} \left\{ \exp \left(\alpha \left(i \frac{k\pi}{b-a}, \tau \right) + \beta \left(i \frac{k\pi}{b-a}, \tau \right)' \bar{V} - i \frac{k\pi a}{b-a} \right) \beta \left(i \frac{k\pi}{b-a}, \tau \right)' \right\} \omega_k \right\} \end{aligned}$$

$$\begin{aligned} &\frac{\partial \mu_2^{\mathbb{Q}^+}(t, \tau)}{\partial V_t} \Big|_{V_t = \bar{V}} \\ &= \frac{1}{2} \frac{\partial^2 \beta'(u, \tau)}{\partial u^2} \Big|_{u=0} + \left(\frac{\partial \alpha(u, \tau)}{\partial u} \Big|_{u=0} + \frac{\partial \beta'(u, \tau)}{\partial u} \Big|_{u=0} \bar{V} \right) \frac{\partial \beta'(u, \tau)}{\partial u} \Big|_{u=0} \\ &\quad - \frac{1}{2} \left\{ \sum_{k=1}^{N-1} \text{Re} \left\{ \exp \left(\alpha \left(i \frac{k\pi}{b-a}, \tau \right) + \beta \left(i \frac{k\pi}{b-a}, \tau \right)' \bar{V} - i \frac{k\pi a}{b-a} \right) \beta \left(i \frac{k\pi}{b-a}, \tau \right)' \right\} \omega_k \right\}. \end{aligned}$$

A.9 Physical Expected Quadratic Payoff and the Drift

In the main paper, we can see from Eq. (13) that the difference between the physical expected quadratic payoff $\mathbb{E}_t[r_{t,t+\tau}^2]$ and the physical expected realized variance $\sigma_{t,\tau}^2$ is due to the non-zero drift ($\mu_{t,\tau} \neq 0$). If the drift equals zero, the expected quadratic payoff is exactly the same as the expected realized variance. In Figure A7, we plot the average term structure of the drift ($\mu_{t,\tau}$) and the squared drift ($\mu_{t,\tau}^2$). As expected, the term structure of the drift is essentially flat. However, the term structure of squared drift is increasing over the 1-month to 12-month horizon. The fact that the squared drift is higher at longer investment horizons explains the much larger discrepancy between the expected quadratic payoff and the expected realized variance at longer horizons. In order to investigate the fluctuations in the drift and the squared drift, we report the descriptive statistics of both quantities in Table A4. We find that both the drift and the squared drift of any investment horizon have significant variations over time.

A.10 Normality Assumption and Physical Expected Quadratic Loss and Gain

To derive Eq. (13) in the main paper, we assume that log returns are normally distributed. This means that the wedge between the physical expected quadratic gain and loss is a function of $\mu_{t,\tau}$. In general, this wedge could also be driven by other factors, e.g., skewness. In this section, we relax the normality assumption and instead assume that the log return $r_{t,t+\tau}$ follows a binormal distribution. This is an analytically tractable distribution that accommodates empirically plausible values of skewness and kurtosis and nests the familiar Gaussian distribution.³ The density function of the binormal distribution with parameters (m, σ_1, σ_2) is given by:

$$f(x) = A \exp\left(-\frac{1}{2}\left(\frac{x-m}{\sigma_1}\right)^2\right) I(x < m) + A \exp\left(-\frac{1}{2}\left(\frac{x-m}{\sigma_2}\right)^2\right) I(x \geq m), \quad (\text{A.17})$$

where $A = \frac{1}{\sigma_1 + \sigma_2} \sqrt{\frac{2}{\pi}}$. We notice that m is the mode, and up to a multiplicative constant, σ_1^2 and σ_2^2 are interpreted as variances of returns, conditional on returns being less than the mode,

³The binormal distribution was first introduced by Gibbons and Mylroie (1973). Further, see Bangert et al. (1986), Kimber and Jeynes (1987), and Toth and Szentimrey (1990), among others, for examples of using the binormal distribution in data modeling, statistical analysis and robustness studies.

and conditional on returns being greater than the mode, respectively. Specifically, we have:

$$\text{Var}[r \mid r < m] = \left(1 - \frac{2}{\pi}\right) \sigma_1^2 \quad \text{and} \quad \text{Var}[r \mid r \geq m] = \left(1 - \frac{2}{\pi}\right) \sigma_2^2. \quad (\text{A.18})$$

This property that disentangles the upside and the downside variance with respect to the mode is considered to be the most important characteristic of the binormal distribution. Clearly, if $\sigma_2 = \sigma_1$, then the binormal distribution is reduced to a normal distribution with mean m and variance σ_1^2 .

The binormal distribution can alternatively be parameterized by the mean μ , the variance σ^2 , and the Pearson mode skewness p . The parametrization (μ, σ, p) is related to the original parametrization (m, σ_1, σ_2) through the following mapping:

$$\begin{cases} \sigma^2 &= \left(1 - \frac{2}{\pi}\right) (\sigma_2 - \sigma_1)^2 + \sigma_1 \sigma_2 \\ p &= \frac{\sigma_2 - \sigma_1}{\sigma} \sqrt{\frac{2}{\pi}} \end{cases} \quad \text{and} \quad \mu = m + \sigma p. \quad (\text{A.19})$$

Likewise, the original parametrization (m, σ_1, σ_2) is related to the alternative parametrization (μ, σ, p) through the following mapping:

$$m = \mu - \sigma p \quad \text{and} \quad \begin{cases} \sigma_1 &= \sigma \left(-\sqrt{\frac{\pi}{8}} p + \sqrt{1 - \left(\frac{3\pi}{8} - 1\right) p^2} \right) \\ \sigma_2 &= \sigma \left(\sqrt{\frac{\pi}{8}} p + \sqrt{1 - \left(\frac{3\pi}{8} - 1\right) p^2} \right). \end{cases} \quad (\text{A.20})$$

Finally, the Pearson mode skewness p and the traditional skewness s of the binormal distribution are related by the following equation:

$$s = p \left(1 - (\pi - 3) p^2\right). \quad (\text{A.21})$$

This means that, with an estimate of the traditional skewness, the Pearson mode skewness can be found as a root of the following depressed cubic equation:

$$x^3 - \frac{1}{\pi - 3} x + \frac{s}{\pi - 3} = 0. \quad (\text{A.22})$$

The roots of a depressed cubic equation can easily be obtained from the Cardano formula.

Assuming that log returns are binormally distributed, their moments as well as those of their loss and gain components can easily be obtained via the moment-generating function $M(u) = \mathbb{E}[\exp(ur)]$ and the truncated moment-generating function $M(u; x) = \mathbb{E}[\exp(ur) I(r \geq x)]$. These functions are explicitly given by:

$$M(u) = \frac{2\sigma_1}{\sigma_1 + \sigma_2} \exp\left(mu + \frac{\sigma_1^2 u^2}{2}\right) \Phi(-\sigma_1 u) + \frac{2\sigma_2}{\sigma_1 + \sigma_2} \exp\left(mu + \frac{\sigma_2^2 u^2}{2}\right) \Phi(\sigma_2 u) \quad (\text{A.23})$$

and

$$M(u; x) = \begin{cases} M(u) - \frac{2\sigma_1}{\sigma_1 + \sigma_2} \exp\left(mu + \frac{\sigma_1^2 u^2}{2}\right) \Phi\left(\frac{x-m}{\sigma_1} - \sigma_1 u\right) & \text{if } x < m \\ \frac{2\sigma_2}{\sigma_1 + \sigma_2} \exp\left(mu + \frac{\sigma_2^2 u^2}{2}\right) \Phi\left(-\frac{x-m}{\sigma_2} + \sigma_2 u\right) & \text{if } x \geq m, \end{cases} \quad (\text{A.24})$$

where Φ is the standard normal cumulative distribution function.

The first- and second-order derivatives of the function $M(u)$ with respect to u , and evaluated at $u = 0$, are the expected return and the expected squared return. This yields

$$\begin{aligned} \mathbb{E}[r] &= m + (\sigma_2 - \sigma_1) \sqrt{\frac{2}{\pi}} = \mu \\ \mathbb{E}[r^2] &= m^2 + \sigma_1^2 - \sigma_1 \sigma_2 + \sigma_2^2 + 2m(\sigma_2 - \sigma_1) \sqrt{\frac{2}{\pi}} = \mu^2 + \sigma^2. \end{aligned} \quad (\text{A.25})$$

Likewise, the first- and second-order derivatives of the function $M(u; x)$ with respect to u , and evaluated at $u = 0$ and $x = 0$, are the expected gain and the expected squared gain. We have the following:

$$\mathbb{E}[g] = \begin{cases} \frac{2\sigma_1}{\sigma_1 + \sigma_2} \left[m \left(\Phi\left(\frac{m}{\sigma_1}\right) + \frac{\sigma_2 - \sigma_1}{2\sigma_1} \right) + \sigma_1 \left(\phi\left(\frac{m}{\sigma_1}\right) + \frac{\sigma_2^2 - \sigma_1^2}{2\sigma_1^2} \sqrt{\frac{2}{\pi}} \right) \right] & \text{if } 0 < m \\ \frac{2\sigma_2}{\sigma_1 + \sigma_2} \left[m \Phi\left(\frac{m}{\sigma_2}\right) + \sigma_2 \phi\left(\frac{m}{\sigma_2}\right) \right] & 0 \geq m, \end{cases} \quad (\text{A.26})$$

and

$$\mathbb{E}[g^2] = \begin{cases} \frac{2\sigma_1}{\sigma_1 + \sigma_2} \left[(m^2 + \sigma_1^2) \left(\Phi\left(\frac{m}{\sigma_1}\right) + \frac{\sigma_2(m^2 + \sigma_2^2) - \sigma_1(m^2 + \sigma_1^2)}{2\sigma_1(m^2 + \sigma_1^2)} \right) \right. \\ \quad \left. + m\sigma_1 \left(\phi\left(\frac{m}{\sigma_1}\right) + \frac{\sigma_2^2 - \sigma_1^2}{\sigma_1^2} \sqrt{\frac{2}{\pi}} \right) \right] & \text{if } 0 < m \\ \frac{2\sigma_2}{\sigma_1 + \sigma_2} \left[(m^2 + \sigma_2^2) \Phi\left(\frac{m}{\sigma_2}\right) + m\sigma_2 \phi\left(\frac{m}{\sigma_2}\right) \right] & 0 \geq m. \end{cases} \quad (\text{A.27})$$

Given the estimates of the return moments and the gain moments, we can obtain estimates of the loss moments using the identities $l = g - r$ and $l^2 = r^2 - g^2$; thus,

$$\mathbb{E}[l] = \mathbb{E}[g] - \mathbb{E}[r] \quad \text{and} \quad \mathbb{E}[l^2] = \mathbb{E}[r^2] - \mathbb{E}[g^2]. \quad (\text{A.28})$$

Utilizing equations A.27 and A.28, we estimate the physical expected quadratic loss and gain. Then we correlate these quantities to the ones estimated under the normality assumption in Eq. (13) of the main paper. In Table A5, we report the correlations between the physical expected quadratic loss (gain) estimated under the assumption of normal and binormal distributed log returns from 1 to 12 months. We find that correlations of these two physical expected quadratic loss (gain) range between 0.984 (0.980) and 0.996 (0.993). We also plot these correlations for all maturities in Figure A8.

In Table A6, we report the mean of physical expected quadratic loss (gain) estimated under the assumption of normal and binormal distributed log returns over the 1- to 12-month horizon. Similar to the term structures of the expected quadratic loss (gain) under the normality assumption, under the more relaxed assumption of binormal distribution, we also find a decreasing (increasing) average term structure for the expected quadratic loss (gain).

A.11 Rolling Window Parameter Estimation

We use the full sample to estimate the drifts in returns and the expectations of realized variances over different periods. These estimated quantities enter Eq. (13) to compute the expectations of the truncated returns. Such an approach may incur forward-looking bias in the parameter estimation

stage. To check the robustness of our results, we estimate drifts and expected realized variances using a set of different rolling windows: 60, 72, 84, 96, 108, and 120 months of daily data. While our sample of intradaily returns starts in January 1990, our sample of risk-neutral quantities starts in January 1996 due to option data availability. This six-year gap in the data makes the use of a 60-month window size efficient. A shorter rolling window is undesirable to study the term structure of expectations of truncated returns as the estimations are much less accurate for longer horizon investments.

Based on the rolling-window estimated drifts and realized variance expectations, we compute the expected quadratic gain and loss. To make comparisons to the full-sample expected quadratic gain and loss, in Figure A9, we plot the correlation between the rolling-window and full-sample expected quadratic gain in blue (loss in black) over the 1-month horizon. We find that all correlations are well above 0.7. Except for the expected quadratic gain using a 60-month window size, all other correlations are higher than 0.77. The correlation can be as high as 0.88 for an expected quadratic loss based on a 120-month rolling window.

To further compare the full-sample and rolling-window expected truncated returns, we fix the rolling window size to 120 months, and for each investment horizon τ , we estimate time-series regressions of the following form:

$$\begin{aligned}\mathbb{E} [g_{FS,\tau}^2] &= \alpha + \beta_{roll}^g \mathbb{E} [g_{roll,\tau}^2] + \varepsilon \\ \mathbb{E} [l_{FS,\tau}^2] &= \alpha + \beta_{roll}^l \mathbb{E} [l_{roll,\tau}^2] + \varepsilon.\end{aligned}\tag{A.29}$$

If these expectations constructed from rolling-window parameters are exactly the same as the ones constructed from full-sample parameters, we should find that β_{roll}^g and β_{roll}^l are not statistically and significantly different from one. In Figure A10, we plot these coefficients for investment horizons from 1 to 12 months. In Panel A, we see that β_{roll}^g s are not statistically different from one over horizons of 6, 7, and 8 months. On the other hand, we find in Panel B that β_{roll}^l is not statistically different from one over a 3-month horizon. Overall, because of the difficulty of predicting returns over short horizons, the quantities constructed from the full-sample parameters inevitably differ from the rolling-window counterparts.

A.12 Transition Equation for the Three Factors in the Andersen et al. (2015b) Model

The transition equation for the three factors in the AFT model is:

$$V_{t+1} = \Phi_0 + \Phi_1 V_t + \Sigma(V_t)^{1/2} \varepsilon_{t+1}, \quad (\text{A.30})$$

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, $\bar{\lambda}_-^{\mathbb{P}} = 2/(\lambda_-^{\mathbb{P}})^2$, and Δt is set to 1/252 to reflect a daily time step.

Moreover, the conditional covariance matrix of the transition noise is:

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

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

Figure A1: Term Structure of Forward Variance Prices and Quadratic Risk Premium

In this figure, in Panel A we plot the term structure of average prices of forward variance as defined in Dew-Becker et al. (2017) for the risk-neutral expected quadratic payoff and its loss and gain components. All forward variance prices are reported in annualized volatility terms, $100 \times \sqrt{12 \times F_t^n}$. In Panel B, we plot the mean S&P 500 quadratic risk premium and its loss and gain components. The sample period is from January 1996 to December 2015.




Figure A2: Level and Slope of Quadratic Risk Premium Term Structure

In this figure, in Panel A, we plot the level (6-month maturity) of the term structure of the S&P 500 quadratic risk premium and its loss and gain components. In Panel B, we plot the slope (12-month minus 2-month maturity) of the term structure of these quantities. The sample period is from January 1996 to December 2015.





Figure A3: Factors Estimates Across Models

In this figure we plot the filtered factors across models ($\eta = 0$ vs $\eta \neq 0$). The sample period is from January 1996 to December 2015.




Figure A4: Average Risk-Neutral Moments

In this figure we plot the observed and model-implied average term structure of risk-neutral moments. The sample period is from January 1996 to December 2015.

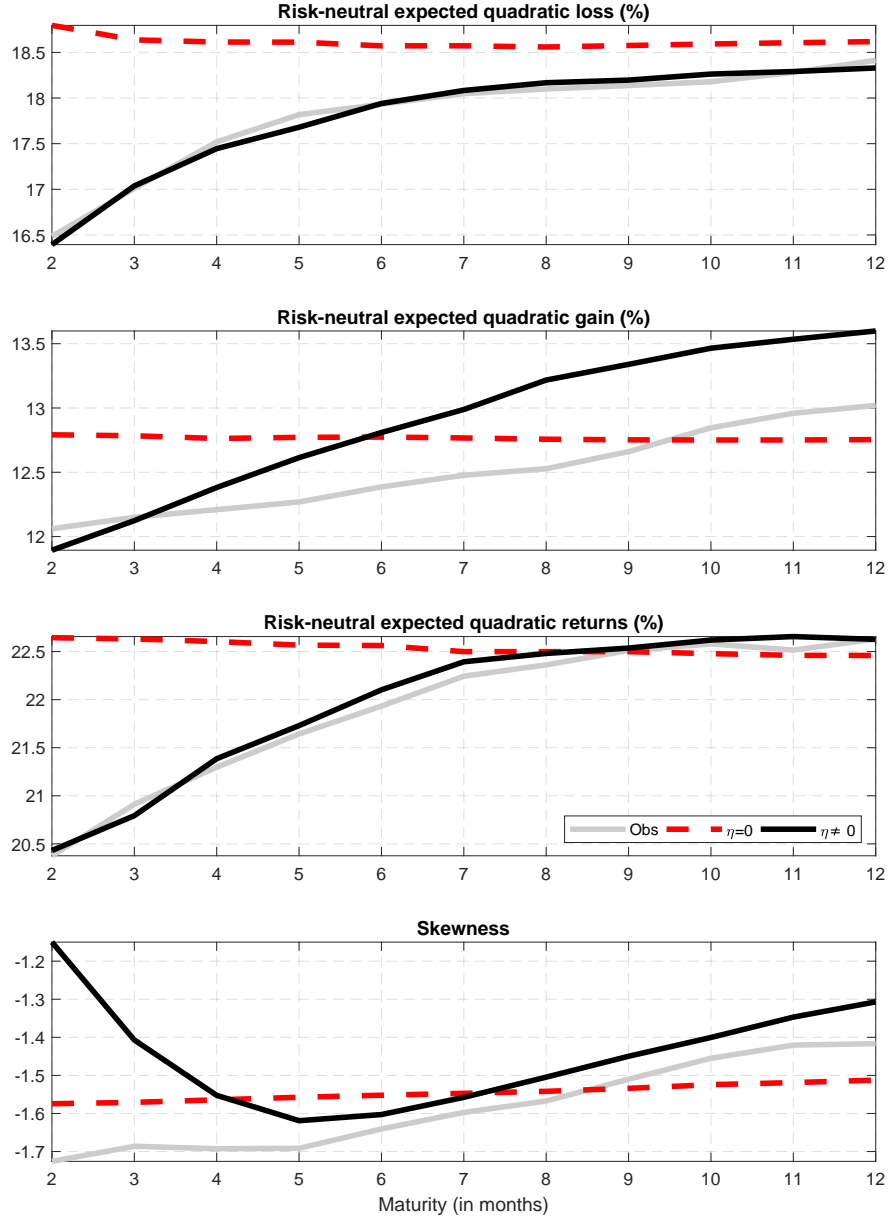


Figure A5: Factors Estimates: Kalman Filter vs Extended Kalman Filter

Focusing on the AFT4 model, in this figure we plot the filtered factors across two estimation methods: the classic Kalman Filter (KF) and the extended Kalman Filter (E-KF). The sample period is from January 1996 to December 2015.

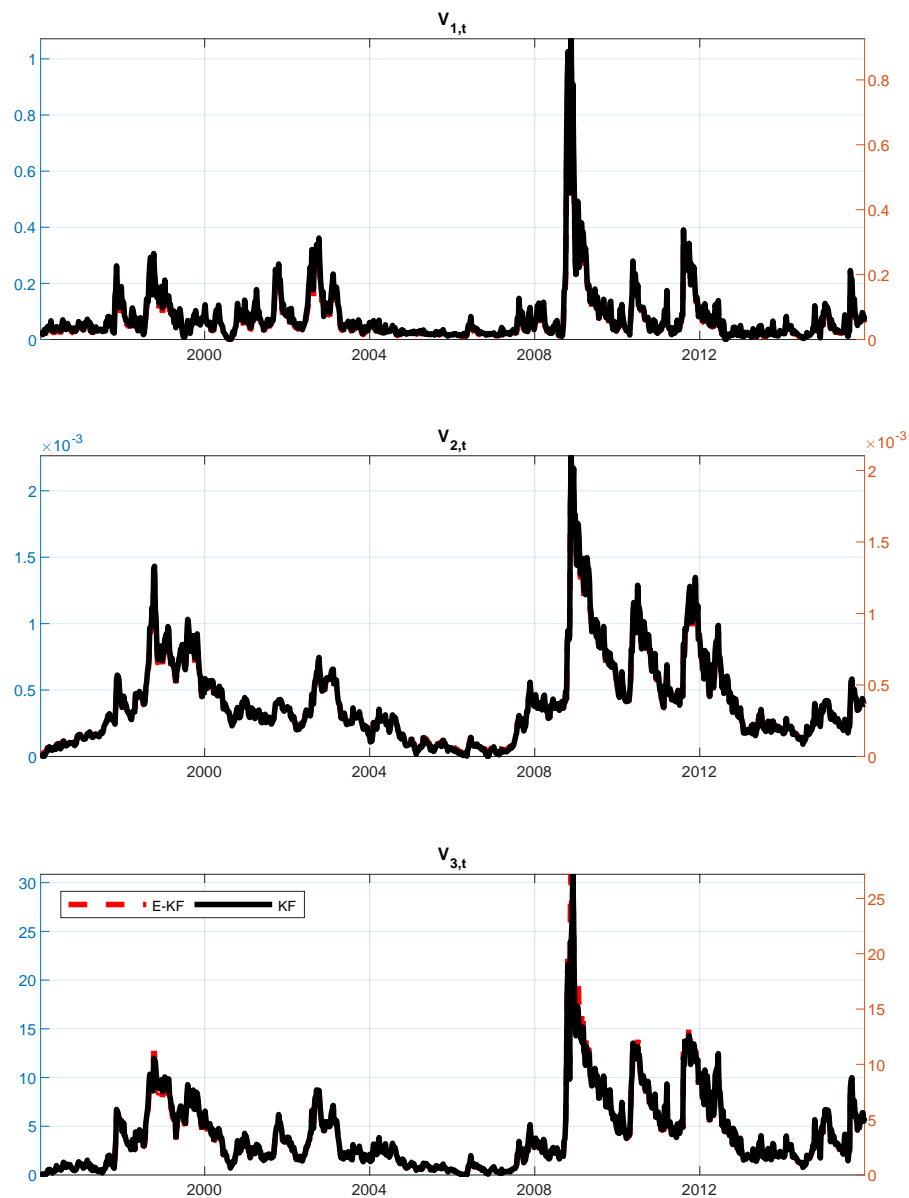


Figure A6: Average Risk-Neutral Moments: Kalman Filter vs Extended Kalman Filter

In this figure we plot the observed and AFT implied average term structure of risk-neutral moments for two estimation methods: the classic Kalman Filter (KF) and the extended Kalman Filter (E-KF). The sample period is from January 1996 to December 2015.




Figure A7: Term Structure of the Drift

In this figure, in Panels A and B, we plot the average term structure of the drift (μ) and the squared drift (μ^2), respectively. All reported values are annualized, and μ is in percentage units, while μ^2 is in squared percentage units. The sample period is from January 1996 to December 2015.

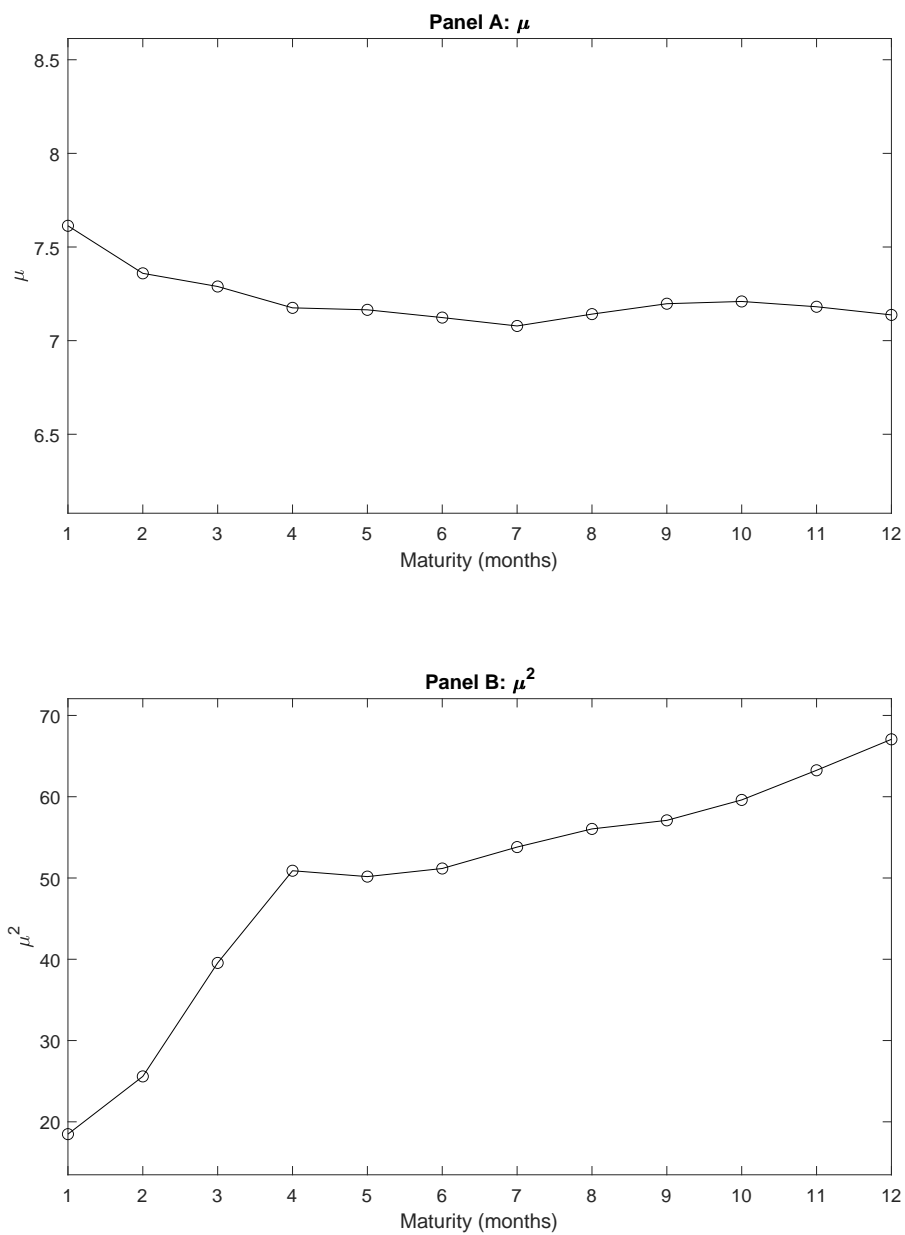


Figure A8: Correlation between Normal and Binormal Expected Quadratic Gain or Loss

In this figure, we plot the correlation between physical expected quadratic gain (loss) estimated with the assumption of normally distributed log returns and physical expected quadratic gain (loss) estimated with the assumption of binormally distributed log returns. We plot this correlation for 1 to 12 months' maturity. The sample period is from January 1996 to December 2015.




Figure A9: Correlation between Full Sample and Rolling Window Expected Quadratic Gain or Loss

In this figure, we plot the correlation between 1-month-horizon physical expected quadratic gain (loss) constructed using parameters estimated by full sample and physical expected quadratic gain (loss) constructed using parameters estimated by rolling window. The parameters are estimated using a set of different rolling windows: 60, 72, 84, 96, 108 or 120 months of daily data, respectively. The sample period is from January 1996 to December 2015.




Figure A10: Full Sample versus Rolling Window Expected Quadratic Gain or Loss

In this figure, in Panels A and B we plot the coefficients β_{roll}^g and β_{roll}^l from Eq. (A.29), respectively. We plot these quantities for maturities from 1 to 12 months and include a 95% confidence band for the coefficients. The rolling window quantities are constructed using parameters estimated with a rolling window of 120 months of daily data. The sample period is from January 1996 to December 2015.

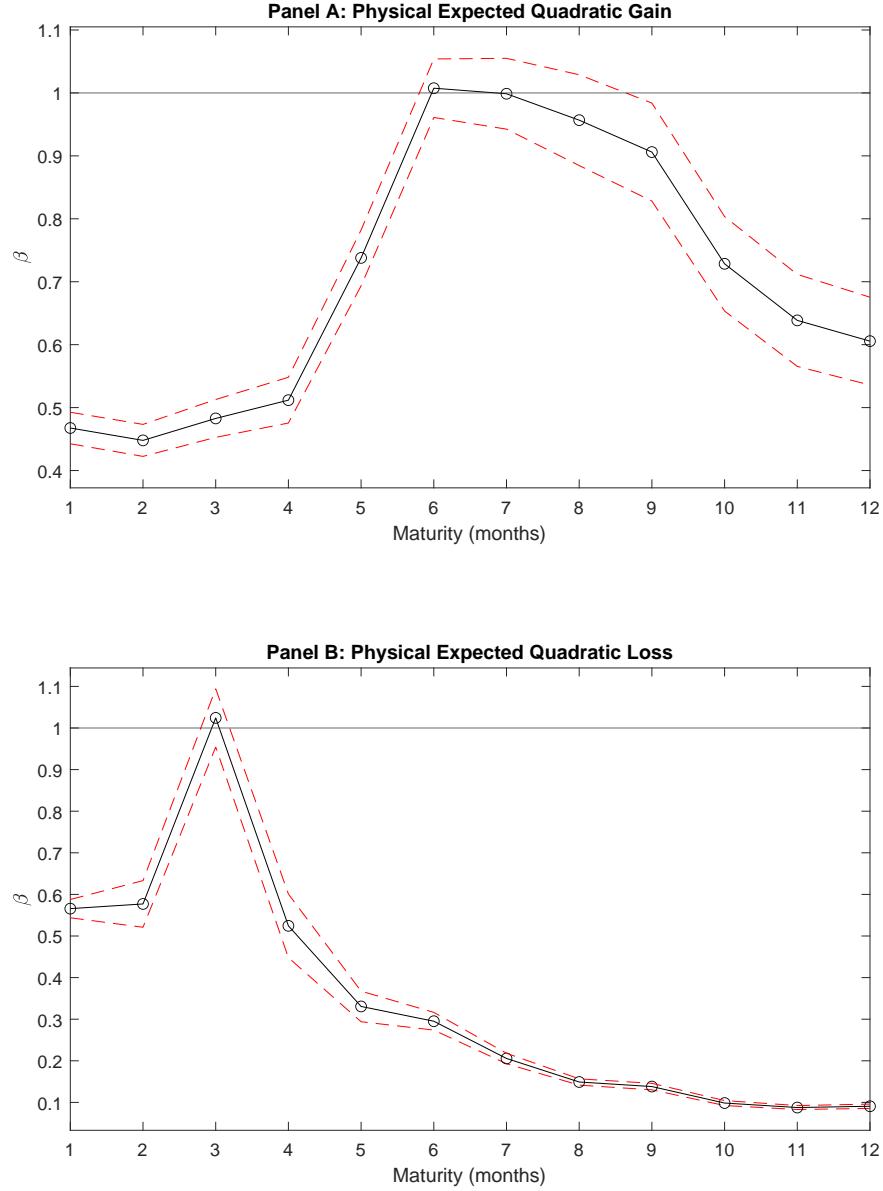


Table A1: $\eta = 0$ vs $\eta \neq 0$: Risk-Neutral Moments RMSEs

In this table we report the root mean squared error

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

where Mom_t^{Mkt} is the time t risk-neutral moment value observed on the market, and 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		Volatility		Skewness	
	$\eta = 0$	$\eta \neq 0$	$\eta = 0$	$\eta \neq 0$	$\eta = 0$	$\eta \neq 0$	$\eta = 0$	$\eta \neq 0$
2	1.49	0.68	1.41	0.82	1.87	0.31	0.19	0.47
3	1.20	0.50	1.34	0.63	1.44	0.11	0.18	0.27
4	0.85	0.44	1.26	0.69	1.12	0.18	0.15	0.19
5	0.61	0.46	1.32	0.80	0.89	0.21	0.15	0.13
6	0.46	0.54	1.34	1.00	0.63	0.21	0.13	0.09
7	0.54	0.66	1.39	1.19	0.44	0.18	0.11	0.09
8	0.79	0.83	1.45	1.37	0.28	0.15	0.12	0.11
9	1.02	0.95	1.53	1.51	0.14	0.12	0.12	0.15
10	1.23	1.03	1.59	1.63	0.17	0.16	0.13	0.17
11	1.23	1.11	1.60	1.72	0.28	0.20	0.14	0.19
12	1.25	1.22	1.65	1.80	0.39	0.28	0.15	0.22
Avg	0.97	0.76	1.44	1.20	0.70	0.19	0.14	0.19

Table A2: **Risk-Neutral Parameter Estimate**

In this table we report risk-neutral parameter estimates (Est.) and their respective standard deviation (Std.) for the five different versions of the Andersen et al. (2015a) model that we consider in the main paper.

	AFT0		AFT1		AFT2		AFT3		AFT4	
Par.	Est.	Std.	Est.	Std.	Est.	Std.	Est.	Std.	Est.	Std.
ρ_1	0.99	0	0.99	0	-0.99	3.6e-4	-0.98	0.17	-0.99	1.1e-6
\bar{v}_1	9.0e-8	1.2e-12	1.1e-4	2.6e-5	1.6e-3	1.8e-4	3.2e-6	6.7e-7	1.3e-3	8.0e-3
κ_1	27.26	1.6e-4	0.39	0.04	1.28	0.07	12.14	2.5e-3	7.67	0.12
σ_1	7.2e-5	6.9e-10	3.7e-3	4.8e-3	0.02	8.0e-3	1.3e-3	6.2e-7	0.25	1.3e-3
ρ_2	-0.99	4.0e-9	0.99	4.1e-10	*	*	-0.99	1.5e-16	0.99	0
\bar{v}_2	0.44	6.9e-6	0.14	0.02	*	*	4.2e-4	2.4e-5	9.4e-5	7.9e-5
κ_2	0.09	5.9e-9	1.2e-3	1.4e-3	*	*	2.6e-5	5.5e-8	0.46	0.04
σ_2	0.02	1.7e-7	0.06	2.0e-3	*	*	1.5e-4	4.2e-6	9.3e-3	3.8e-3
μ_3	*	*	*	*	280.13	69.97	1.78	0.07	4.17	3.38
κ_3	*	*	*	*	4.11	0.14	1.25	6.3e-4	3.86	0.11
ρ_3	*	*	*	*	0.99	0	0.58	1.1e-3	0.58	0.05
c_0^+	*	*	7.8e-4	1.2e-4	2.2e-8	4.1e-4	0.26	9.3e-3	0.24	0.11
c_1^-	*	*	21421.25	447.95	364.65	29.86	0.04	7.9e-4	0.05	6.2e-3
c_1^+	*	*	25.81	6.04	1.03	3.05	0.04	7.9e-4	2.1e-8	1.0e-3
c_2^-	*	*	299.10	0.15	*	*	1851.16	100.87	2091.80	853.32
c_2^+	*	*	5.92	0.19	*	*	1851.16	100.87	28.71	0.03
c_3^-	*	*	*	*	0.43	0.17	2.05	0.03	2.37	0.99
λ_-	*	*	87.26	3.12	24.82	0.60	37.91	0.67	55.49	26.64
λ_+	*	*	1.51	5.7e-4	1472.55	557.60	37.91	0.67	44.04	13.94
μ_1	*	*	2.7e-6	1.7e-5	0.03	0.04	1.8e-3	3.2e-5	0.03	0.03
c_0^-	*	*	0.29	0.55	0.04	0.03	0.26	9.3e-3	0.06	0.03
c_3^+	*	*	*	*	72.76	18.20	2.05	0.03	2.36	0.95
η	*	*	*	*	-0.26	0.03	-0.17	1.8e-3	-0.16	1.0e-3

Table A3: **Kalman vs Extended Kalman: Risk-Neutral Moments RMSEs**

Focusing on the AFT4 model, this table compares two estimation techniques: the Kalman (KF) and the Extended Kalman (E-KF), by reporting the root mean squared error

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

where Mom_t^{Mkt} is the time t risk-neutral moment value observed on the market, and 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		Volatility		Skewness	
	E-KF	KF	E-KF	KF	E-KF	KF	E-KF	KF
2	0.65	0.68	0.82	0.82	0.31	0.31	0.49	0.47
3	0.48	0.50	0.62	0.63	0.11	0.11	0.28	0.27
4	0.44	0.44	0.67	0.69	0.17	0.18	0.19	0.19
5	0.44	0.46	0.80	0.80	0.20	0.21	0.14	0.13
6	0.52	0.54	0.97	1.00	0.20	0.21	0.09	0.09
7	0.62	0.66	1.14	1.19	0.17	0.18	0.09	0.09
8	0.76	0.83	1.35	1.37	0.14	0.15	0.11	0.11
9	0.87	0.95	1.49	1.51	0.12	0.12	0.14	0.15
10	0.97	1.03	1.59	1.63	0.15	0.16	0.16	0.17
11	1.04	1.11	1.69	1.72	0.19	0.20	0.19	0.19
12	1.12	1.22	1.77	1.80	0.27	0.28	0.22	0.22
Avg.	0.72	0.76	1.17	1.20	0.18	0.19	0.19	0.19

Table A4: **Descriptive Statistics for the Drift**

In this table, we report a set descriptive statistics for the drift (μ) and the squared drift (μ^2), including mean, median and standard deviation. These statistics are all based on annualized values and are reported in percentage units for μ , and squared percentage units for μ^2 for 1 to 12 months' maturity. Data are from January 1996 to December 2015.

Maturity	μ			μ^2		
	Mean	Median	Std. Dev.	Mean	Median	Std. Dev.
1	7.61	6.32	12.81	18.48	3.77	83.51
2	7.36	6.23	9.97	25.59	7.41	86.68
3	7.29	6.36	10.26	39.55	11.39	162.73
4	7.18	6.11	10.06	50.89	14.29	178.14
5	7.16	6.14	8.32	50.17	17.84	109.65
6	7.12	6.01	7.19	51.17	20.13	91.74
7	7.08	5.99	6.49	53.80	22.70	91.49
8	7.14	6.18	5.75	56.04	26.72	88.73
9	7.20	6.38	4.93	57.09	31.21	77.53
10	7.21	6.45	4.42	59.61	35.13	74.74
11	7.18	6.48	4.18	63.25	38.88	76.23
12	7.14	6.47	4.02	67.06	42.27	77.41

Table A5: Correlations between Binormal and Normal Quadratic Payoff

In this table, we report the correlation between the physical expected quadratic gain (loss) estimated with the assumption of normally distributed log returns in Eq. (13) of the main paper, and the physical expected quadratic gain (loss) estimated using the assumption of binormally distributed log returns in Eq. (A.27). We report these correlations for 1 to 12 months' maturity. Data are from January 1996 to December 2015.

Maturity	Corr. Expected Quadratic Gain	Corr. Expected Quadratic Loss
1	0.984	0.994
2	0.982	0.995
3	0.980	0.995
4	0.984	0.996
5	0.986	0.996
6	0.986	0.988
7	0.987	0.989
8	0.988	0.989
9	0.989	0.988
10	0.991	0.989
11	0.990	0.984
12	0.993	0.993

Table A6: Term Structure of Binormal and Normal Quadratic Payoff

In this table, we report the mean of the physical expected quadratic gain (loss) estimated with the assumption of normally distributed log returns in Eq. (13) of the main paper, and the mean for the physical expected quadratic gain (loss) estimated using the assumption of binormally distributed log returns in Eq. (A.27) for 1 to 12 months' maturity. All reported values are monthly and in squared percentage units. Data are from January 1996 to December 2015.

Maturity	Norm. Quadratic Gain	Binorm. Quadratic Gain	Norm. Quadratic Loss	Binorm. Quadratic Loss
1	15.42	13.09	10.88	13.22
2	16.47	13.92	10.06	12.61
3	17.39	15.09	10.06	12.36
4	18.28	15.97	9.96	12.26
5	18.91	16.30	9.10	11.70
6	19.55	16.94	8.35	10.96
7	20.09	17.44	7.91	10.56
8	20.65	17.92	7.43	10.17
9	21.14	18.38	6.95	9.71
10	21.62	18.91	6.57	9.27
11	22.08	19.43	6.32	8.98
12	22.52	19.77	6.14	8.89

References

- Andersen, T., Fusari, N. and Todorov, V. (2015a). Parametric inference and dynamic state recovery from option panels, *Econometrica* **83**(3): 1081–1145.
- Andersen, T., Fusari, N. and Todorov, V. (2015b). The Risk Premia Embedded in Index Options, *Journal of Financial Economics* **117**(3): 558–584.
- Bakshi, G., Kapadia, N. and Madan, D. (2003). Stock return characteristics, skew laws and the differential pricing of individual equity options, *Review of Financial Studies* **16**(1): 101–143.
- Bangert, U., Goodhew, P. J., Jeynes, C. and Wilson, I. H. (1986). Low Energy (2-5 keV) Argon Damage in Silicon, *Journal of Physics D: Applied Physics* **19**: 589–603.
- Barndorff-Nielsen, O. E., Kinnebrock, S. and Shephard, N. (2010). “Measuring downside risk: realised semivariance,” in *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*, eds. T. Bollerslev, J. Russell, and M. Watson, Oxford: Oxford University Press, pp. 117–136.
- Bollerslev, T., Tauchen, G. and Zhou, H. (2009). Expected stock returns and variance risk premia, *Review of Financial Studies* **22**(11): 4463–4492.
- Christoffersen, P., Jacobs, K. and Ornathanalai, C. (2012). Dynamic jump intensities and risk premiums: Evidence from S&P500 returns and options, *Journal of Financial Economics* **106**(3): 447–472.
URL: <http://www.sciencedirect.com/science/article/pii/S0304405X12001122>
- Dew-Becker, I., Giglio, S., Le, A. and Rodriguez, M. (2017). The Price of Variance Risk, *Journal of Financial Economics* **123**(2): 225–250.
- Fang, F. and Oosterlee, C. (2008). A novel pricing method for european options based on Fourier-cosine series expansions, *SIAM J. Scientific Computing* **31**: 826–848.
- Feunou, B., Jahan-Parvar, M. R. and Okou, C. (2017). Downside Variance Risk Premium, *Journal of Financial Econometrics* **16**(3): 341–383.
URL: <https://dx.doi.org/10.1093/jjfinec/nbx020>
- Feunou, B., Lopez Aliouchkin, R., Tédongap, R. and Xu, L. (2019). Loss uncertainty, gain uncertainty, and expected stock returns, *Working Paper, Bank of Canada, Syracuse University and ESSEC Business School*.
- Gibbons, J. F. and Myroie, S. (1973). Estimation of Impurity Profiles in Ion-Implanted Amorphous Targets Using Joined Half-Gaussian Distributions, *Applied Physics Letters* **22**: 568–569.
- Kilic, M. and Shaliastovich, I. (2019). Good and bad variance premia and expected returns, *Management Science* **65**(6): 2522–2544.
URL: <https://doi.org/10.1287/mnsc.2017.2890>
- Kimber, A. C. and Jeynes, C. (1987). An Application of the Truncated Two-Piece Normal Distribution to the Measurement of Depths of Arsenic Implants in Silicon, *Journal of the Royal Statistical Society: Series B* **36**: 352–357.
- 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, P. and Yan, S. (2010). Crashes, volatility, and the equity premium: Lessons from S&P 500 options, *The Review of Economics and Statistics* **92**(2): 435–451.
URL: <http://www.jstor.org/stable/27867547>
- Toth, Z. and Szentimrey, T. (1990). The Binormal Distribution: A Distribution for Representing Asymmetrical but Normal-Like Weather Elements, *Journal of Climate* **3**(1): 128–137.