Investor Irrationality and the Nasdaq Bubble*

Gurdip Bakshi[a†] Liuren Wu[b‡]

[a]*Smith School of Business, University of Maryland, College Park, MD 20742, USA*

[b]*Zicklin School of Business, Baruch College, One Bernard Baruch Way, New York, NY 10010, USA*

Abstract

We exploit the information in the options market to study the risk and risk premium variations around the Nasdaq bubble period. In particular, we investigate whether the dramatic rise and fall of the Nasdaq can be justified by changes in return risk, or there were corresponding unusual shifts in how investors price risks. We specify a model that accommodates fluctuations in both risk levels and market prices of different sources of risks, and we estimate the model using time-series returns and option prices on the Nasdaq 100 index. Our analysis reveals three key pieces of evidence that foretell the arrival and burst of a bubble. First, during the Nasdaq bubble period, return volatility increased together with the rising index level, even though the two tend to move in opposite directions in general. Second, while the market price of volatility risk is strongly negative historically as investors dislike high volatility levels and volatility risk, the market price of volatility risk declined in absolute magnitude and approached zero at the end of 1999. In contrast to the average risk premium at normal market conditions, aversion to volatility risk completely disappeared during the height of the bubble period. Third, the options market showed an increasing market price of downside jump risk that peaked three month prior to the bursting of the bubble, highlighting the increasing demand for hedging against the potential crash of the Nasdaq market valuation. On the other hand, the market price of the downside jump risk dropped suddenly three months prior to the burst, pointing to capital flight from the Nasdaq market and the unwinding of the associated option hedging positions. Our approach suggests that having options on the bubble asset provides a window into investor behavior unavailable from earlier bubbles.

JEL Classification: G10, G12, G13, C51.

Keywords: Nasdaq; bubble; risks; risk premiums; market prices of risks; variations in risk preferences; diffusion risk; volatility risk; downside jump risk; options.

---

*This research was funded by BSI Gamma Foundation. We thank Andrew Ang, Doron Avramov, Peter Carr, Mike Chernov, Josh Coval, Bjorn Eraker, Soeren Hvidkjaer, Michael Johannes, Nengjiu Ju, Alok Kumar, Pete Kyle, Stephen Leroy, Mark Loewenstein, Dilip Madan, Jun Pan, Neil Pearson, George Pennacchi, Allen Poteshman, Hersh Shefrin, Andrei Simonov, Rene Stulz, Hao Zhou, and seminar participants at the 2006 Western Finance Association meetings in Colorado, the BSI Gamma Foundation conference in Frankfurt, and the University of Illinois at Urbana-Champaign for comments. Any remaining errors are ours.

†Tel.: +1-301-405-2261; fax: +1-301-405-0359. *E-mail address:* gbakshi@rhsmith.umd.edu.

‡Tel.: +1-646-312-3509; fax: +1-646-312-3451. *E-mail address:* liuren_wu@baruch.cuny.edu.

1. Introduction

The Nasdaq market endured an unusual transformation at the turn of the 21st century. During the one-year period from the end of March 1999 to March 2000, the Nasdaq 100 index, which constitutes the vast majority of the Nasdaq market capitalization, rose by 128%. Stock prices started to decline after that, and by March 2001, the level of the Nasdaq 100 index was about 30% of its peak value in 2000. This phenomenon was labeled by many as the "Nasdaq bubble." Despite the progress made in understanding the consequences and causes of bubbles,[1] the debate continues: (i) What are the defining characteristics of a bubble? (ii) What are the economic forces behind the formation and collapse of stock price bubbles? (iii) How can we predict a bubble *ex ante* or identify a bubble *ex post*?

A bubble often refers to a period of sharp price increase of an asset followed by its dramatic drop, all the while without the corresponding identifiable movement in the underlying fundamentals or cash flows. See, for example, Abreu and Brunnermeier (2003), Allen and Gorton (1993), Gilles and Leroy (1992), Kindleberger (1978), Leroy (2004), Santos and Woodford (1997), and references therein. In the absence of cash flow variations, the price change is presumably due to variations in the pricing kernel, which can be thought of as a reduced-form summary of the market's risk preferences and subjective probability beliefs, e.g., Harrison and Kreps (1979), Scheinkman and Xiong (2003), and Kogan, Ross, Wang, and Westerfield (2006). Thus, it is natural to think that a bubble and its ensuing burst are associated with a dramatic fluctuation in the market prices of various sources of risks on the market. In this sense, identifying a price bubble amounts to identifying systematic variations in the market prices of various sources of risks.

Ever since the classic works by Garber (1979) and Kindleberger (1978), bubbles and bursts have been a recurring theme of the financial market research. In contrast to the famous earlier bubbles, e.g., Dutch tulip mania and South Sea bubble, the Nasdaq bubble segment incorporates a unique additional information source in the form of options quotes on the bubble asset. Each

[1] See, for instance, Abreu and Brunnermeier (2003), Battalio and Schultz (2006), Brunnermeier and Nagel (2004), Dhar and Goetzmann (2005), Griffin, Harris, and Topaloglu (2004), Hong, Scheinkman, and Xiong (2006), Lakonishok, Lee, Pearson, and Poteshman (2006), Leroy (2004), Ofek and Richardson (2002, 2003), Pástor and Veronesi (2005, 2006), and Temin and Voth (2004).

day, option prices at different strikes and maturities provide us with a picture of the risk-adjusted return distribution at different future realizations and over different investment horizons. Thus, by analyzing the variations of the option prices and the underlying index, we can learn how risks and risk premiums have varied during the bubble period.

In this paper, we infer the behaviors of the market prices of various sources of risks in the Nasdaq market by exploiting the information in the options market. We specify a model that accommodates four distinct sources of risks: (i) diffusion return risk, (ii) return volatility risk, (iii) upside-jump risk, and (iv) downside-jump risk. By assigning separate market prices for each of the risk sources, we empirically determine their average magnitudes and their fluctuations around the Nasdaq bubble period using information in the time-series of index returns and option prices.

Through model construction and estimation, we investigate the following hypotheses. First, if stable market price of risk estimates are accompanied by changes in the level of risk that comoves with the rise and the fall of the Nasdaq, we attribute the Nasdaq bubble to corresponding variations in the underlying risk level. On the other hand, if the market prices of different sources of risks underwent dramatic changes around the bubble segment, we argue in favor of either or both of the following two possible scenarios: (i) Investors' risk attitudes changed dramatically around this period, declining risk aversions prior to the bursting of the bubble, and then reverting back to the normal level thereafter. (ii) Investors' subjective probability distribution deviates substantially from the objective distribution. The latter could occur, for instance, if investors believe that the future cash flows to the Nasdaq stocks have far higher growth rates and/or lower risks than dictated by their true distribution. Without distinguishing between the two explanations, we call this outcome as irrational exuberance in the sense popularized by Alan Greenspan, Shiller (2000), Shleifer (2000), and Shefrin (2005). Finally, if neither the ex-ante risk level nor the market prices of risks experienced unusual movements, we classify the dramatic rise and fall of the Nasdaq as a random, albeit extreme, realization.

Our estimation results show that return volatility on the Nasdaq 100 index started to climb up in late 1999, similar to the observation in Schwert (2002). We also find that volatility spiked after the bubble burst, not before. The return volatility calmed down during a short period in the middle of 2000 as the Nasdaq market experienced a short-lived recovery, but volatility rose again as the market started to drop again after that. Thus, we do not regard the rising volatility feature as an explanation for the soaring Nasdaq valuation, but rather we regard its positive co-movement with the market valuation as the first sign of anomalous market conditions: At normal times, it has been well-documented that equity valuation and its volatility tend to move in opposite directions.[2]

On the other hand, our estimation has identified significant variations in the market prices of different sources of risks around the bubble period. At normal market condition, the market price of volatility risk is on average highly negative, showing investors' general dislike for high volatility and volatility risk.[3] However, the market price of volatility risk became close to zero at the end of 1999. Investors' aversion to volatility risk completely disappeared at the height of the Nasdaq valuation. After the bubble burst, investor's dislike for volatility risk reverted to its traditional values. The absolute decline in the market price of volatility risk during the bubble period and its reversal thereafter can be interpreted as evidence on the existence of investor irrationality.

Furthermore, during the three to six months before the burst of the bubble, the market prices of downside jump risk reached historical high levels, showing the increasing demand for hedging against potential crashes in the Nasdaq market valuation. Intriguing, the market price of the downside jump risk dropped suddenly three months prior to the burst of the bubble. The sudden drop points to capital flight from the Nasdaq market and hence the unwinding of the associated option hedging positions.

Overall, we have identified anomalous variations in both the risk level and the market prices of both diffusion risk and downside jump risk during the bubble period. These variations are all interrelated in generating the Nasdaq bubble and its ensuing burst. The declining market price of volatility risk reflects investor irrationality and the result of increasing speculative buying pressure. Such speculative buying activities, which possibly interacted with the rising return volatility, also generated increasing demand for downside protection, thereby raising premiums for out-of-the-money put op-

[2]Black (1976) uses the leverage effect to explain the empirical regularity: With fixed debt, declining equity value increases the firm's financial leverage and hence increases the riskiness of the equity. Various other explanations have also been proposed in the economics literature, e.g., French, with G. William Schwert, and Stambaugh (1987), Haugen, Talmor, and Torous (1991), Campbell and Hentschel (1992), Campbell and Kyle (1993), and Bekaert and Wu (2000).

[3]Similar stylized evidence has also been found in Bakshi and Kapadia (2003), Carr and Wu (2008), and Bondarenko (2004).

tions. Finally, the sudden drop in the hedging demand suggests fund withdrawals from the Nasdaq market, signaling the ensuing burst of the bubble.

The next section formalizes the model structure that decomposes the sources of risks and characterizes the market price of risks. Section 3 outlines the estimation procedure. Section 4 investigates several hypotheses on the identification and causes of the bubble. Section 5 concludes.

2. The sources of risks and market prices of risks

The aim of the modeling framework is to first decompose the index return into several distinct risk sources and then assign a separate market price to each risk source. The decomposition of risk and market price of risk allows us to investigate whether risk levels and/or market prices of various sources of risks experienced systematic variations during the bubble period. The further decomposition of different risk sources allows us to identify the nature of investor irrationality that fed the Nasdaq bubble.

Formally, we fix a filtered probability space $\{\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P}\}$ and model the index return as a time-changed Lévy process:

$$\ln S_t/S_0 = \mu[.] + W^s_{\mathcal{T}_t} + J^+_{\mathcal{T}_t} + J^-_{\mathcal{T}_t} - \xi \mathcal{T}_t, \tag{1}$$

where S_t is the time-t price level, and $\mu[.]$ represents the expected return that is commensurate with the underlying risks and their market pricing. Equation (1) decomposes market risks into four sources:

- W^s_t denotes a standard Brownian motion capturing the diffusive return risk,
- J^+_t denotes a purely upside jump component,
- J^-_t denotes a purely downside jump component, and,
- $\mathcal{T}_t$ is a stochastic time change applied to return innovations to generate stochastic volatility.

Equation (1) posits a distinct role for diffusion risks, upside jump risks, and downside jump risks, and allows for differential pricing of each component. The term ξ represents the concavity adjustment induced by the three return innovation components, $\xi = k_{W^s}[1] + k_{J^+}[1] + k_{J^-}[1]$, where $k_X[u]$ denotes the cumulant exponent of the corresponding Lévy component $X = W_t^s$, J_t^+, and J_t^-, respectively,

$$k_X[u] \equiv \frac{1}{t}\ln \mathbb{E}\left(e^{uX_t}\right), \qquad u \in \mathcal{D} \subset \mathbb{C}. \tag{2}$$

Cumulant exponents are traditionally defined on the positive real half-line, but for convenience we extend the definition to the subset of the complex plane where the expectation is well-defined.

The cumulant exponent of the Brownian motion component is $k_{W^s}[u] = \frac{1}{2}u^2$. The cumulant exponent of the two jump components depend on the specification of the jump structure. We model the up and down jumps via their respective Lévy densities $(\pi^+[x], \pi^-[x])$, which control the arrival rates of return jumps of size x. For tractability, we let the arrival rates follow exponentially dampened power laws:

$$\pi^+[x] = \begin{cases} \lambda e^{-\beta_+ x} x^{-2}, & x > 0, \\ 0, & x < 0. \end{cases}, \quad \pi^-[x] = \begin{cases} 0, & x > 0, \\ \lambda e^{-\beta_-|x|}|x|^{-2}, & x < 0. \end{cases}, \tag{3}$$

with $(\lambda, \beta_+, \beta_-) \in \mathbb{R}^+$. The power term $|x|^{-2}$ generates infinite number of jumps as the jump size approaches zero, a property that has gained empirical support (see Carr, Geman, Madan, and Yor (2002), Wu (2006), and Bakshi, Carr, and Wu (2008)). The exponential terms in equation (3) dampen the tails so that all return moments are finite, and the different dampening coefficients β_+ and β_- allow the up and down jumps to generate different thickness for the right and left tails of the index return, respectively. The scaling coefficient λ determines the overall level of the jump frequency and hence their overall contribution to return volatility and higher return moments. This jump specification provides flexibility in accommodating negatively-skewed and fat-tailed risk-neutral return distributions observed in the options market. Furthermore, compared to the classic compound Poisson jump model of Merton (1976), the jump specification in (3) presents a natural separation for downside and upside jumps, which we allow investors to price differently.

Given the Lévy density specification, the cumulant exponent of the jump components can be derived via the Lévy-Khintchine theorem (Sato (1999)):

$$k_J[u] = \int_{\mathbb{R}^0} \left(e^{ux} - 1 - ux\,1_{|x|<1}\right) \pi[x]\, dx. \tag{4}$$

Under the Lévy densities in (3), the cumulant exponents for the upside and downside jump components become (Wu (2006)),

$$k_{J^+}[u] = \lambda(\beta_+ - u)\ln(1 - u/\beta_+) + uC_+, \qquad k_{J^-}[u] = \lambda(\beta_- + u)\ln(1 + u/\beta_-) + uC_-, \tag{5}$$

where C_+ and C_- are two constants that depend on the model parameters and the truncation function and will eventually be canceled out with the corresponding concavity adjustment.

To generate stochastic volatility, we apply a stochastic time change $\mathcal{T}_t$ to the three return components. We assume that the stochastic time change is locally predictable and continuous, and specify $\mathcal{T}_t$ via the instantaneous volatility rate V_t:

$$\mathcal{T}_t \equiv \int_0^t V_u\, du. \tag{6}$$

Under this time change, the instantaneous return variance per unit calendar time is $V_t\left(1 + \lambda\left(\beta_+^{-1} + \beta_-^{-1}\right)\right)$, which is proportional to V_t. Hence V_t generates stochastic return volatility, and the volatility rate summarizes the variation in both the instantaneous variance of the Brownian motion component and the arrival rates of up and down jumps. We model the volatility rate V_t dynamics under measure $\mathbb{P}$ by a square-root process,

$$dV_t = \kappa(\theta - V_t)\, dt + \omega\sqrt{V_t}\, dW_t^v, \tag{7}$$

where κ controls the mean-reversion speed, θ captures the long-run mean, and W_t^v denotes standard Brownian motion that is correlated with the Brownian motion in the index return: $\rho dt \equiv \mathbb{E}(dW_t^s dW_t^v)$. The dynamics of V_t can be thought of as capturing the variation of the overall risk in the Nasdaq market.

To link the above statistical dynamics to market prices, we specify the following pricing kernel, which assigns a separate market price for each source of risk:

$$\mathcal{M}_t = \exp(-rt)\exp\left(-\gamma_s W^s_{\mathcal{T}_t} - \gamma_+ J^+_{\mathcal{T}_t} - \gamma_- J^-_{\mathcal{T}_t} - \gamma_v W^v_{\mathcal{T}_t} - \hbar\mathcal{T}_t\right), \tag{8}$$

where r is the continuously compounded spot interest rate of the relevant maturity, which we assume deterministic, and the second term is an exponential martingale that defines the measure change from the statistical measure $\mathbb{P}$ to the risk-neutral measure $\mathbb{Q}$, with $\hbar \equiv \frac{1}{2}\gamma_s^2 + k_{J^+}[-\gamma_+] + k_{J^-}[-\gamma_-] + \frac{1}{2}\gamma_v^2 + \gamma_s\gamma_v\rho$ being the convexity adjustment.

The innovation in specification (8) is that it assigns separate market prices for the diffusive return risk (W^s_t), the upside jump risk $\left(J^+_t\right)$, the downside jump risk $\left(J^-_t\right)$, and the stochastic volatility risk (W^v_t) via $(\gamma_s, \gamma_+, \gamma_-, \gamma_v)$, respectively. These market prices capture the differences between the risk-neutral dynamics and the statistical dynamics, as a result of (i) investors' risk aversions to different sources of risks and/or (ii) deviations of investors' subjective beliefs from statistical realities. We identify the market prices without making a distinction between the two explanations.

Under the pricing kernel in (8), the return dynamics under the risk-neutral measure $\mathbb{Q}$ become,

$$\ln S_t/S_0 \;=\; rt + W^{\mathbb{Q}}_{\mathcal{T}_t} + J^{+\mathbb{Q}}_{\mathcal{T}_t} + J^{-\mathbb{Q}}_{\mathcal{T}_t} - \xi^{\mathbb{Q}}\mathcal{T}_t. \tag{9}$$

with the concavity adjustment under measure $\mathbb{Q}$ given analogously by $\xi^{\mathbb{Q}} = \frac{1}{2} + k^{\mathbb{Q}}_{J^+}[1] + k^{\mathbb{Q}}_{J^-}[1]$. Under the risk-neutral measure $\mathbb{Q}$, the Lévy densities of the two jump components are exponentially-tilted versions of their physical counterparts:

$$\pi^{+\mathbb{Q}}[x] \;=\; \begin{cases} \lambda e^{-(\beta_+ + \gamma_+)x} x^{-2}, & x > 0, \\ 0, & x < 0. \end{cases}, \quad \pi^{-\mathbb{Q}}[x] = \begin{cases} 0, & x > 0, \\ \lambda e^{-(\beta_- - \gamma_-)|x|} |x|^{-2}, & x < 0. \end{cases} \tag{10}$$

Thus, the Lévy densities under $\mathbb{P}$ and $\mathbb{Q}$ share the same dampened power law functional form, but with different exponential dampening coefficients $\beta^{\mathbb{Q}}_+ = \beta_+ + \gamma_+$ and $\beta^{\mathbb{Q}}_- = \beta_- - \gamma_-$. Positive market prices on the two jump risks increase the dampening on upside jumps and reduce the dampening on downside jumps. As a result, the return innovation distribution becomes more negatively skewed under the risk-neutral measure $\mathbb{Q}$.

Under the risk-neutral measure, the volatility rate dynamics become,

$$dV_t = \kappa^{\mathbb{Q}} \left(\theta^{\mathbb{Q}} - V_t \right) dt + \omega \sqrt{V_t}\, dW_t^v, \tag{11}$$

where the risk-neutral mean reversion speed becomes $\kappa^{\mathbb{Q}} = \kappa + \gamma_v \omega + \gamma_s \omega \rho$ and the risk-neutral long-run mean becomes $\theta^{\mathbb{Q}} = \kappa\theta / \kappa^{\mathbb{Q}}$. A negative market price of diffusive volatility risk (γ_v) or a positive market price of diffusive return risk (γ_s) together with negative correlation between the return and volatility reduces the mean reversion speed and increases the long-run mean of the volatility rate under the risk-neutral measure.

The following theorem decomposes the return risk premium into three components:

Theorem 1 *Suppose that (i) the asset dynamics is of the class (1), (ii) the jump structure is of the exponentially dampened power law class (3), (iii) the stochastic time change is of the class (6), and (iv) the pricing kernel specification is of the class (8). Then the instantaneous return risk premium is ηV_t, where V_t measures the time-t risk level and η measures the risk premium per unit risk, which can be decomposed into three components:*

$$\eta \equiv \underbrace{\eta_d}_{\textit{Market Price of Diffusion}} + \underbrace{\eta_{J^+}}_{\textit{Market Price of Upside Jumps}} + \underbrace{\eta_{J^-}}_{\textit{Market Price of Downside Jumps}} \tag{12}$$

with the market price on each source of risk determined by its respective cumulant exponent:

$$\begin{aligned} \eta_d &= \gamma_s + \gamma_v \rho, \\ \eta_{J^+} &= k_{J^+}[1] - k_{J^+}[1-\gamma_+] + k_{J^+}[-\gamma_+], \\ \eta_{J^-} &= k_{J^-}[1] - k_{J^-}[1-\gamma_-] + k_{J^-}[-\gamma_-], \end{aligned} \tag{13}$$

where $k_{J^-}[u]$ and $k_{J^-}[u]$ are presented in (5).

Proof: See Appendix A.

Consistent with economic theory (e.g., Harrison and Kreps (1979) and Cox and Huang (1989)), the risk premium per unit risk from each Lévy risk component reflects the difference between the cumulant exponents under $\mathbb{P}$ and $\mathbb{Q}$, evaluated at $u = 1$. When jump risk is not priced and volatility is constant, the risk premium expression reduces to the counterpart in Merton (1976). Otherwise, the market price of the diffusive volatility risk γ_v also contributes to the diffusion risk premium η_d through the correlation between the Brownian motion in the return (W^s) and the Brownian motion in the volatility rate (W^v). Furthermore, the market prices of upside jump risk (γ_+) and downside jump risk (γ_-) generate two additional risk premiums, η_{J^+} and η_{J^-}.

Unlike the traditional focus on the behavior of total return risk premium in some early studies of bubbles and bursts, the decomposition in (12) is imperative to our analysis. The first question to ask is whether a risk premium change is due to a change in return risk (V_t) or a change in its market price (the risk premium per unit return risk, η). Our explicit distinction and separate estimation of return risk and market price of risk can be useful for addressing whether the bubble is due to variation in risk or variation in market price per unit risk. Furthermore, it has been recognized that the financial market is not one dimensional and that several risk sources are priced separately. By classifying the risk sources and assigning a different market price of risk to each risk source, we can go one step further in asking which source of risk and/or its market price variation was responsible for the Nasdaq bubble.

3. The estimation approach

Guided by the implications of Theorem 1, we estimate the risk dynamics (V_t) jointly with the market prices of risks based on the time series of returns and option prices on the Nasdaq 100 index tracking stock. We select a four year sample period from March 17, 1999 to February 19, 2003 to coincide with the rise and fall of Nasdaq. Option observations on the Nasdaq 100 tracking stock are taken from the"Ivy DB" data set sold by OptionMetrics. For model estimation, we sample the option data weekly on every Wednesday. The procedure generates 20,160 option prices over 206 weeks. The weekly returns on the index tracking stock are computed based on Wednesday closing prices.

To describe the estimation procedure, we define y_t as out-of-the-money option prices scaled by Black-Scholes vega at time t, and $O[V_t;\Theta]$ as the model-generated value as a function of the unobservable volatility rate, V_t, and the model parameters set Θ. The scaling of the option price by the vega is meant to control for moneyness and maturity effects in the estimation (e.g., Bakshi, Carr, and Wu (2008)).

In order to compute the model values for the options, we first derive in analytical form the generalized Fourier transform of the log index return under the risk-neutral measure $\mathbb{Q}$, the details of which are shown in Appendix B. Then, we compute the option values via fast Fourier inversion of the generalized transform as in Carr and Wu (2004).

The approach to estimate the model parameters Θ and to extract the volatility rates V_t at each date t is based on a four-step iterative procedure:

Step I. Given initial guess on parameters Θ, we cast the model into a state-space form and extract the unobservable volatility rates $\{V_t : t = 1,\ldots,T\}$ using an extended version of the Kalman filter. Relying on an Euler approximation of the volatility dynamics in (7), the state propagation equation is:

$$V_t \;=\; A + \Phi V_{t-1} + \sqrt{\mathcal{G}_{t-1}}\,\varepsilon_t, \qquad V_t \in \mathbb{R}^+ \tag{14}$$

where ε_t is an i.i.d. standard normal innovation, $A = (1-\exp(-\kappa\Delta t))\theta$, $\Phi = \exp(-\kappa\Delta t)$, $\mathcal{G}_{t-1} = \omega^2 V_{t-1}\Delta t$ with $\Delta t = 1/52$ denoting the weekly time interval used in our estimation. The measurement equations are based on out-of-the-money option prices, assuming additive, normally-distributed measurement errors:

$$y_t \;=\; O[V_t;\Theta] + e_t. \tag{15}$$

With the vega weighting on the option prices, we assume that the pricing errors are i.i.d. normal with zero mean and constant volatility. The dimension of the measurement equation varies over time as we have different number of options available at different dates. In this setup, the state-propagation equations are Gaussian linear, but the measurement equations is nonlinear in the state variable. We use an extended version of the Kalman filter, the unscented Kalman filter (UKF), to handle the nonlinearity in the measurement equation, the details of which are provided in Appendix C.

Step II. Given the unscented Kalman filter time-$(t-1)$ forecasts of option prices, $\overline{y}_t$, and its conditional covariance matrix $\overline{\Omega}_t$, we build the log likelihood from options under the assumption of normally distributed forecasting errors,

$$l_t^O[\Theta] = -\frac{1}{2}\ln\left|\overline{\Omega}_t\right| - \frac{1}{2}\left((y_t-\overline{y}_t)^\top\left(\overline{\Omega}_t\right)^{-1}(y_t-\overline{y}_t)\right). \tag{16}$$

Step III. Having extracted the return volatility rate $\widehat{V}_{t-1}$ from index options at time $(t-1)$, we compute the statistical density of the log index return, $\ln(S_t/S_{t-1})$, as a function of $\widehat{V}_{t-1}$ by applying fast Fourier inversion to the return characteristic function under the statistical measure $\mathbb{P}$, which we derive in analytical form in Appendix D. Taking logarithm of the density values yields the conditional log likelihood on the index return, $l_t^s[\Theta]$.

Step IV. In the final step, we choose model parameters to maximize the joint conditional log likelihoods from both options and index returns,

$$\max_\Theta \mathcal{L}[\Theta] \equiv \sum_{t=1}^{T}\left(l_t^O[\Theta]+l_t^s[\Theta]\right), \tag{17}$$

where we assume conditional independence between the options forecasting errors and the index returns. We iterate through the four steps until the desired level of convergence is achieved.

According to equation (12) in Theorem 1, we decompose the return risk premium per unit risk into three components: (i) the diffusion risk premium, which is jointly determined by γ_s and γ_v, (ii) the upside jump risk premium, which is determined by γ_+, and (ii) the downside jump risk premium, which is determined by γ_-. To enhance identification, we demean the return series by regressing log index prices against time and set γ_s to zero. Then, we identify the diffusion risk premium through the estimation of γ_v, the market price of the diffusive volatility risk.

Furthermore, since our objective is to analyze whether market prices of different risks have experienced unusual variations during the bubble period, we also estimate an extended version of the model, where we allow the market prices to vary over time. Specifically, we use $z_t \equiv (\gamma_v, \ln\beta_+^{\mathbb{Q}}, \ln\beta_-^{\mathbb{Q}})$

to denote the vector that we allow to vary over time, with the following auxiliary dynamics:

$$z_t = \theta_z(1-\phi) + \phi z_{t-1} + \sigma_x x_t, \tag{18}$$

where θ_z denotes the long-run mean, the scaler ϕ captures the mean-reverting tendency of z_t, and x_t embodies time-t information relevant for updating market prices of risks. We normalize x_t to have zero mean and identity covariance matrix, with σ_x being a constant scalar. With (β_+, β_-) fixed, time variations in $(\beta_+^{\mathbb{Q}}, \beta_-^{\mathbb{Q}})$ reflect variations in the market prices of up and down jumps (γ_+, γ_-) because $\beta_+^{\mathbb{Q}} = \beta_+ + \gamma_+$ and $\beta_-^{\mathbb{Q}} = \beta_-^{\mathbb{Q}} - \gamma_-$. The log transformation of $\beta_+^{\mathbb{Q}}$ and $\beta_-^{\mathbb{Q}}$ in the state vector z_t expands the domain of the state vector to the whole real line.

Given assumption (18), we re-estimate the model with time-varying market prices using the same procedure, but with an expanded state vector $\mathcal{V}_t' = (V_t, z_t)$. To avoid the complication of convexity terms for option pricing, we take the market prices as deterministically time-varying and treat equation (18) as a conditional forecasting equation analogous to the GARCH specification for volatilities (Engle (1982) and Bollerslev (1986)). Hence, the filtered updates on z_t can be regarded as time-t forecasts of future market prices of risks.

4. What caused the Nasdaq bubble?

To study the underlying causes behind the unusual transformation of the Nasdaq market, we pose the following testable hypotheses.

- First, if stable market price estimates are accompanied by changes in the volatility rate level that move with the rise and fall of the Nasdaq, we trace the bubble to unusual variations in the underlying volatility level of the Nasdaq market.
- Second, if the market prices of risks show unusual variations around the bubble, we argue in favor of either or both of the following two possible scenarios: (i) Investors' risk attitudes changed around this period: Declining risk aversions prior to the bursting of the bubble, and then reverting back to the normal level thereafter. (ii) Investors' subjective probability distribu-

tion deviates substantially from the objective distribution (Kogan, Ross, Wang, and Westerfield (2006)). The latter could occur, for instance, if the investor believes that the future cash flow to the Nasdaq stocks have far higher growth rates and/or lower risks than dictated by their true distribution, e.g., Ofek and Richardson (2003). Without distinguishing between the two explanations, we call declining market prices of risks as irrational exuberance.

Finally, if neither the ex-ante volatility rate level nor the market prices of risks experienced unusual movements, we classify the dramatic rise and fall of the Nasdaq as a random realization.

4.1. The return volatility increased with the soaring Nasdaq valuation

To investigate the first hypothesis, we examine the time-variation in return volatility V_t during the Nasdaq bubble period. Table 1 presents the model parameter estimates when the market prices of risks are held constant. All the risk dynamics parameters are estimated with high statistical significance, suggesting that our empirical procedure is successful in identifying the statistical properties of volatility. Specifically, the large variance of variance coefficient $\omega = 0.5603$ implies that return volatility itself is highly volatile. The large estimate for $\kappa = 6.2608$ suggests quick mean-reversion in Nasdaq volatility. The long-run mean (θ) estimate of 0.1949 represents an average statistical return volatility level of 47.12%, computed as $\sqrt{\theta(1+\lambda(\beta_+^{-1}+\beta_-^{-1})}$, which is several folds larger than the corresponding estimate from the S&P 500 index (e.g., Andersen, Benzoni, and Lund (2002)). The correlation estimate $\rho = -0.8878$ supports the economically plausible result that W_t^s and W_t^v are strongly inversely related.

Table 2 reports the estimation results with time-varying market prices. The parameter estimates for the return volatility dynamics are similar to those reported in Table 1 for the constant market price of risk case, showing the robustness of estimation and model specification.

To assess how return volatility varied around the Nasdaq bubble, we plot in Figure 1 the time series of the filtered volatility rate $\widehat{V}_t$. The left panel is the time series extracted from the model with time-varying market prices. The right panel is the time series from the constant market price model. In both panels, the solid lines represent the time series of the volatility rate V_t, and the dashed lines denote the fluctuation of the Nasdaq 100 index tracking stock level from March 1999 to April 2001. The time series in the two panels show similar patterns. Hence, both market price of risk specifications yield the same implication for the relation between the return volatility and the Nasdaq valuation.

[Figure 1 about here.]

During the two-year period, the Nasdaq 100 index tracking stock rose from about $50 to about $117 in March 2000 and ending at about $33 in April 2001. The return volatility rate started at a relatively low level in March 1999, but steadily increased as the Nasdaq 100 index climbed up. One way of evaluating how much volatility rose is to note that the estimated Nasdad tracking stock return volatility on March 31, 1999 was 29.2% versus 47.2% on March 29, 2000. The positive co-movement between the index level and return volatility prior to the burst of the bubble is not a common phenomenon in the stock index market. In contrast, at normal market conditions, it is well-documented that stock level and stock volatility tend to move in opposite directions. Therefore, we do not regard the rising return volatility in late 1999 as a normal market behavior or a rationale for the soaring Nasdaq valuation, but rather regard it as an aberration from normal market conditions and a sign foretelling the arrival and burst of a potential bubble.

The return volatility continued rising as the Nasdaq started to fall and reached 62.3% at its peak on April 19, 2000. The return volatility generally remained high after the bursting of the bubble in March 2000. Return volatility subsided for a few months when the Nasdaq stabilized, but trended upwards thereafter as the Nasdaq index continued its downfall. Thus, the bursting of the bubble seems to have triggered a negative market sentiment that kept the Nasdaq index volatility high.

The rising Nasdaq volatility during the bubble period has also been found by Schwert (2002). In the speculative behavior model of Harrison and Kreps (1978), investors bought the Nasdaq so as to sell it to others for more than they think it is worth. What we find here shows that if investors did speculate this way, they did so in the presence of pronounced return volatility. As Nasdaq valuation continued to rise, the market became increasingly uncertain about its sustainability. The increasing return volatility reflects the heightened market uncertainty and speculation about the direction of the market.

Recently, Scheinkman and Xiong (2003) propose an equilibrium model where overconfidence generates disagreements among agents regarding asset fundamentals. This disagreements, combined with short-sale constraints, push up both the asset prices and asset volatility. On the other hand, Pástor and Veronesi (2006) argue that a positive correlation is possible when volatility is largely idiosyncratic. According to this interpretation, Nasdaq stocks would have experienced abnormally high idiosyncratic volatilities as the bubble built up. Furthermore, the fact that volatility remained high after the burst of the bubble suggests an increase in the systematic volatility component for Nasdaq stocks, consistent with the prediction of Pástor and Veronesi (2005) during technology revolutions. Although the interpretations are different, both theories suggest that a combination of rising volatility with rising market valuation is likely to be followed by a subsequent downturn, due to either the burst of a bubble in case of the overconfidence theory or the transformation of idiosyncratic risks at the beginning of a technology revolution into systematic risks as the technology matures into the mainstream industry.

4.2. Diffusion risk premium approached zero as the bubble built up

The remainder of the paper addresses two questions. First, is the market price of diffusion risk and jump risk time-varying? If so, what role does time-varying market price of risks play in explaining the behavior of Nasdaq after controlling for the risk level?

Table 2 reports the parameter estimates accounting for time-variation in market prices of risks. The three parameters $(\gamma_v, \beta_+^{\mathbb{Q}}, \beta_-^{\mathbb{Q}})$ in the third panel of Table 2 represent their respective long-run means (θ_z) as defined in (18).

At the crux of the time-varying market price of risk hypothesis is the estimate for σ_x^2, which is reported in the group on the right side of Table 2. The estimate is highly significant with a t-statistic of 26.05. The magnitude of 0.0939 implies an instantaneous volatility (σ_x) of 30.64%, suggesting that the market price variation is not only statistically significant but also large in magnitude. Thus, the null hypothesis of constant market prices of risks is strongly rejected.

The left panel in Figure 2 plots the time series of the market price of the diffusive volatility risk,

γ_v, as the solid line and the Nasdaq index as the dashed line. The market price of volatility risk started negative and stayed negative on average, suggesting that investors do not like exposures to high volatility level and high variations in return volatility. A series of recent studies have also identified average negative market prices on the market volatility risk, e.g., Bakshi and Kapadia (2003), Bondarenko (2004), and Carr and Wu (2008), suggesting that negative market price of volatility risk is a robust empirical regularity. However, in late 1999, accompanying the soaring Nasdaq valuation, the market price of volatility risk spiked up and even became positive momentarily. This unusual movement suggests that investors were no longer averse to volatility risk during this period.

Soon after, the market price of volatility risk reversed its course to become more negative, and reached historical lows ensuing the burst of the bubble, demonstrating investors' renewed dislike for volatility risk. For instance, the estimated market price of volatility risk (γ_v) started at -2.04 on March 31, 1999, but became even temporarily positive at 0.25 on November 24, 1999. The market price became negative again soon after and reached historical lows at -3.85 by August 23, 2000.

[Figure 2 about here.]

Given the pervasive finding of a negative correlation between return and volatility innovations ($\rho < 0$), a negative market price of the diffusive volatility risk generates a positive diffusion risk premium on the index return. To illustrate this idea more clearly, the right panel in Figure 2 plots the diffusion risk premium (η_d) generated from exposures to per unit of volatility risk. As seen from (13), this standardized return risk premium is given by $\gamma_v\rho$, with γ_s normalized to zero. The risk premium started positive at slightly below two but became negative in late 1999. Therefore, based on our estimation and after controlling for return risk variations, we do identify a period of extremely low or even negative diffusion return risk premium per unit risk that preceded the rise of the bubble. On the other hand, after the bubble burst, the diffusion risk premium per unit risk became strongly positive again, suggesting a reversion of investors attitudes to risk aversion.

4.3. *Put premiums reached historical highs, but suddenly dropped three months prior to the burst*

Investors hold the market portfolio in aggregate. The demand for hedging against downward market jumps forces $\beta_+^{\mathbb{Q}} > \beta_-^{\mathbb{Q}}$, and hence a negatively skewed risk-neutral return distribution. In the left panel of Figure 3, the solid line plots the time series of the difference between the two dampening coefficients, $\beta_+^{\mathbb{Q}} - \beta_-^{\mathbb{Q}}$, which can be regarded as a measure of excess demand for out-of-the-money put options and a gauge of market concerns for downside jumps. This gauge of market concern started positive and stayed positive most of the time during our sample period, indicating that the left tail of the index return is consistently thicker than the right tail under the risk-neutral measure. In economic terms, the demand for hedging against market crashes exceeds the demand for betting on large upside potentials.

[Figure 3 about here.]

Therefore, on the one hand, in late 1999, investors had strong demands in long Nasdaq stocks. This demand fed into the ensuing market valuations and the subsequent correction. On the other hand, with the Nasdaq soaring, players in the options markets grew increasingly apprehensive about the potentially significant market correction. As such, these investors started to buy out-of-the-money puts to hedge their long position in the stock market, driving the prices of these puts much higher than their call option counterparts.

Interestingly, the out-of-the-money put premium as measured by $\beta_+^{\mathbb{Q}} - \beta_-^{\mathbb{Q}}$ reached its peak in the late 1999 and dropped suddenly about three to six months prior to the burst of the bubble. Brunnermeier and Nagel (2004) find that hedge funds did not correct the misvaluations in the Nasdaq stocks by taking offsetting positions. Rather, they unloaded their positions nearly six months prior to the bursting of the bubble. The sudden drop in the out-of-the-money put premium coincided with the their documented unwinding of hedge fund positions in the Nasdaq market, which led to unwinding of the hedging positions in out-of-the-money puts.

Such evidence is also consistent with the theoretical model of Leland (1980), who shows that in equilibrium out-of-the-money options should be bought by investors who think that the market will go up, and thus have long positions in the market. The view of such investors changed three months prior to the burst of the bubble.

Finally, to explore the relation between the market price of jump risk and options trading activities, we also obtain daily open interest information on the index tracking stock from OptionMetrics. The right panel of Figure 3 plots the relative imbalance in open interest between all call options and all put options. The open interest plot shows increasing imbalance before the burst of the bubble. The open interest on puts becomes increasingly large than the call option counterpart. The imbalance reached 75 percent by March 2000. Only after the burst of the bubble did the imbalance start to dissipate.

The plots in the two panels deviate markedly in their behaviors during the three months prior to the burst of the bubble in March 2000. While the market price imbalance $\beta_+^{\mathbb{Q}} - \beta_-^{\mathbb{Q}}$ had dropped by the end of 1999, the open interest imbalance stayed high right to the burst of the bubble. Our conjecture is that on the one hand, the unwinding of the hedging position reduced the premium for out-of-the-money put options; on the other hand, investors who wanted to take short positions in the Nasdaq market but had short sale constraints started to take put positions to achieve short exposure on the Nasdaq market.

4.4. Putting it all together

Our estimation has identified three anomalous shifts in the Nasdaq market that foretell the arrival and burst of the Nadaq bubble. First, prior to the burst of the bubble, the index return volatility increased with the rising market valuation, whereas the two tend to move in opposite directions under normal market conditions. Second, the diffusion risk premium approached zero as the Nasdaq index valuation increases, a sign of irrational exuberance. Third, the out-of-the-money put options became unusually expensive as compared to out-of-the-money call options, revealing the increasing hedging demand for long Nasdaq market positions. Nevertheless, this hedging demand suddenly dropped three months prior to the burst of bubble, a sign that institutional investors started to unload their long positions and hence their hedges in the Nasdaq market. At the same time, overall open interest in put options increased relative to the overall open interest in call options, suggesting that investors started to take more put positions to gain short exposure on the Nasdaq market.

The presence of options trading on the Nasdaq index dramatically enhances our identification of all three pieces of evidence. The options market is the natural market for volatility trading. Thus, option prices make return volatility almost an observable quantity. Furthermore, by combining the information in option prices across strikes and maturities with the time series returns on the index, we have achieved a sharper separation of different sources of risks and the market prices on each risk source. As a result, we are able to observe the diminishing risk premium on small, diffusive risks but increasing risk premium on large, downside movements. It is the combination of the two, together with the rising volatility, that predicts the arrival and burst of a bubble.

5. Conclusions

This paper develops a theoretical and empirical framework to study risk and market price of risks over the Nasdaq bubble period. Our model estimation identifies three interrelated pieces of evidence that reveal unusual variations in both the volatility level and the market prices of diffusion and jump risks prior to the burst of the bubble.

First, return volatility on the Nasdaq 100 index went up with the Nasdaq valuation in late 1999, although at normal times the two often move in the opposite directions.

Second, market prices of both diffusion and jump risks experienced unusual movements around the bubble period. Although investors dislike volatility risk and high volatility at normal times, the aversion to volatility risk declined and became close to zero by the end of 1999.

Third, the market prices of downside jump risk reached historical high levels three months before the bursting of the bubble, reflecting the increasing hedging demand and growing concern for Nasdaq market crashes. This hedging demand dropped suddenly three months prior to the burst of bubble, consistent with earlier evidence that hedge fund unloaded their long positions and hence their hedge in the Nasdaq market three to six months prior to the burst of the bubble. Meanwhile, the open interest in all put option continued to increase relative to open interest in all call options, indicating that investors were taking increasing put options to gain short exposures in the Nasdaq market.

The documented variations in return risk levels, the market price of volatility risk, and the market price of jump risk over the bubble period suggest that the rise and fall of the Nasdaq was not a random aberration. Instead, these variations interact with one another in generating the bubble and its ensuing burst. Our analysis also suggests that options on the bubble asset provide a window into investor behavior that is unavailable from earlier bubbles.

Appendix A. Proof of Theorem 1

With the return dynamics specified as a time-changed Lévy process in (1), we first derive the risk premium induced from each Lévy risk component conditional on a unit risk level $V_t = 1$. The time-t instantaneous return risk premium is then simply the product of the risk premium per unit risk and the risk level V_t.

For each Lévy return risk component, the risk premium per unit risk reflects the difference between the cumulant exponents under $\mathbb{P}$ and $\mathbb{Q}$, evaluated at $u = 1$. Thus, the market prices of the diffusion risk, the upside jump risk, and the downside jump risk are, $\eta_{W^s} \equiv k_{W^s}[1] - k^{\mathbb{Q}}_{W^s}[1]$, $\eta_{J^+} \equiv k_{J^+}[1] - k^{\mathbb{Q}}_{J^+}[1]$, and $\eta_{J^-} \equiv k_{J^-}[1] - k^{\mathbb{Q}}_{J^-}[1]$, respectively.

By the property of exponential martingales that define the measure change in (8), the cumulant exponents of a Lévy risk component under the $\mathbb{P}$ and $\mathbb{Q}$ measures are related by (Küchler and Sørensen (1997)):

$$k^{\mathbb{Q}}_X[u] = k_X[u - \gamma_X] - k_X[-\gamma_X], \tag{A1}$$

where X denotes a Lévy component and γ_X denotes its market price.

For the Brownian motion component, the cumulant exponent under $\mathbb{P}$ is $k_{W^s}[u] = \frac{1}{2}u^2$, and the risk premium per unit risk becomes,

$$\begin{aligned}
\eta_{W^s} &= k_{W^s}[1] - k^{\mathbb{Q}}_{W^s}[1], \\
&= k_{W^s}[1] - \left(k_{W^s}[1 - \gamma_s - \gamma_v\rho] - k_{W^s}[-\gamma_s - \gamma_v\rho]\right) = \gamma_s + \gamma_v\rho,
\end{aligned} \tag{A2}$$

where the market price of volatility risk γ_v contributes to the diffusion return risk premium because of the correlation (ρ) between the volatility risk W^v and diffusion return risk W^s. The market prices of the two jump risk components can be derived by combing (5) and (A1). $\square$

Appendix B. Generalized Fourier transforms and option pricing

Consider the time-t value of an option that matures at $t+\tau$, with the time-to-maturity being τ. To price this option, we first derive the generalized Fourier transform of the log index return, $\ln(S_{t+\tau}/S_t)$, under the risk-neutral measure $\mathbb{Q}$,

$$\phi^{\mathbb{Q}}[u] \equiv \mathbb{E}_t^{\mathbb{Q}}\left(e^{iu(\ln S_{t+\tau}/S_t)}\right) \qquad u \in \mathcal{D} \in \mathbb{C} \tag{B3}$$

where $\mathbb{E}_t^{\mathbb{Q}}$ denotes the expectation operator under measure $\mathbb{Q}$ conditional on time-t filtration $\mathcal{F}_t$. Once we have the generalized Fourier transform, we can apply fast Fourier inversion to the transform to obtain option prices (Carr and Madan (1999) and Carr and Wu (2004)).

Equation (9) writes the risk-neutral return dynamics as a time-changed Lévy process. We can apply the theorem in Carr and Wu (2004) and convert the generalized Fourier transform in equation (B3) into the Laplace transform of the stochastic time change under a new complex valued measure,

$$\phi^{\mathbb{Q}}[u] = \mathbb{E}_t^{\mathbb{Q}}\left(e^{iu(\ln S_{t+\tau}/S_t)}\right) = e^{iur\tau}\, \mathbb{E}_t^{\mathbb{M}}\left(e^{-\psi^{\mathbb{Q}}[u]\mathcal{T}_{t,t+\tau}}\right), \tag{B4}$$

where $\psi^{\mathbb{Q}}[u] = \frac{1}{2}u^2 - k_{J+}^{\mathbb{Q}}[iu] - k_{J-}^{\mathbb{Q}}[iu] + iu\xi^{\mathbb{Q}}$ is the characteristic exponent of the return innovation and $\mathbb{M}$ denotes a new measure. The measure change from $\mathbb{Q}$ to $\mathbb{M}$ is defined by the following exponential martingale:

$$\left.\frac{d\mathbb{M}}{d\mathbb{Q}}\right|_t = \exp\left(iu\left(W_{\mathcal{T}_t}^{s\mathbb{Q}} + J_{\mathcal{T}_t}^{+\mathbb{Q}} + J_{\mathcal{T}_t}^{-\mathbb{Q}}\right) - iu\xi^{\mathbb{Q}}\mathcal{T}_t + \psi^{\mathbb{Q}}[u]\,\mathcal{T}_t\right). \tag{B5}$$

We can drive the volatility rate dynamics (V_t) under the new measure $\mathbb{M}$ according to the Girsanov theorem,

$$dV_t = \left(\kappa\theta - \kappa^{\mathbb{M}}V_t\right)dt + \omega\sqrt{V_t}\,dW_t^{v\mathbb{M}}, \tag{B6}$$

with $\kappa^{\mathbb{M}} = \kappa^{\mathbb{Q}} - iu\omega\rho$. Using (B6) and solving the expectation in (B4), we have:

$$\phi^{\mathbb{Q}}[u] = \exp\left(iur\tau - \frac{2\kappa\theta}{\omega^2} \ln\left(1 - \frac{\vartheta - \kappa^{\mathbb{M}}}{2\vartheta}\left(1 - e^{-\vartheta\tau}\right)\right) - \frac{\kappa\theta}{\omega^2}(\vartheta - \kappa^{\mathbb{M}})\tau - \frac{2\psi^{\mathbb{Q}}\left(1 - e^{-\vartheta\tau}\right)V_t}{2\vartheta - \left(\vartheta - \kappa^{\mathbb{M}}\right)\left(1 - e^{-\vartheta\tau}\right)} \right), \tag{B7}$$

where $\vartheta \equiv \sqrt{(\kappa^{\mathbf{M}})^2 + 2\omega^2\psi^{\mathbb{Q}}}$ and V_t denotes the time-t level of return volatility rate. $\square$

Appendix C. Unscented Kalman filter and maximum likelihood

Kalman filter and its various extensions belong to the state space estimation regime and are based on a pair of state propagation equations and measurement equations. In our application, the state is the volatility rate V_t, which propagates according to the square-root process in equation (7). Its discrete-time version is,

$$V_t = A + \Phi V_{t-1} + \sqrt{\mathcal{G}_{t-1}}\,\varepsilon_t, \tag{C8}$$

which is linear and conditionally normal. In the cases with time-varying market prices, the state vector is $V_t = (V_t, z_t)$, which is expanded to include $z_t = (\ln\beta_+^{\mathbb{Q}}, \ln\beta_-^{\mathbb{Q}}, \gamma_v)$, with equation (18) governing their propagation. The measurement equation, as described in equation (15),

$$y_t = O[V_t;\Theta] + e_t. \tag{C9}$$

is based on the observations on out-of-the-money option prices. After we scale the option values by the Black-Scholes vega, we assume that the pricing errors e_t are independent of the state variable V_t and are i.i.d. normal with zero mean and constant volatility σ_e.

Let $\overline{V}_t, \overline{\Sigma}_t, \overline{y}_t, \overline{\Omega}_t$ denote the time-$(t-1)$ ex ante forecasts of time-t values of the state vector, the covariance of the state, the measurement series, and the covariance of the measurement series. Let $\widehat{V}_t$ and $\widehat{\Sigma}_t$ denote the ex post update, or filtering, on the state vector and its covariance at time t based on observations (y_t) at time t. In the case of linear measurement equation,

$$y_t = HV_t + e_t, \tag{C10}$$

the Kalman filter provides the most efficient updates. The ex ante predictions are

$$\begin{array}{rclrcl} \overline{V}_t & = & \Phi\widehat{V}_{t-1}; & \overline{\Sigma}_t & = & \Phi\widehat{\Sigma}_{t-1}\Phi^\top + \mathcal{G}_{t-1}; \\ \overline{y}_t & = & H\overline{V}_t; & \overline{\Omega}_t & = & H\overline{\Sigma}_t H^\top + \sigma_e^2 I. \end{array} \tag{C11}$$

The ex post filtering updates are

$$\widehat{V}_{t+1} = \overline{V}_{t+1} + K_{t+1}\left(y_{t+1} - \overline{y}_{t+1}\right); \quad \widehat{\Sigma}_{t+1} = \overline{\Sigma}_{t+1} - K_{t+1}\overline{\Omega}_{t+1}K_{t+1}^\top, \tag{C12}$$

where K_{t+1} is the Kalman gain, given by

$$K_{t+1} = \overline{\Sigma}_{t+1}H^\top\left(\overline{\Omega}_{t+1}\right)^{-1}. \tag{C13}$$

In our application, the measurement equation in (C9) is nonlinear. Traditionally, nonlinearity is often handled by the Extended Kalman Filter (EKF), which approximates the nonlinear measurement equation with a linear expansion, evaluated at the predicted states:

$$y_t \approx H\left[\overline{X}_t;\Theta\right]V_t + e_t, \tag{C14}$$

where

$$H\left[\overline{V}_t;\Theta\right] = \left.\frac{\partial O\left[\overline{V}_t;\Theta\right]}{\partial X_t}\right|_{V_t=\overline{V}_t}. \tag{C15}$$

The rest of prediction and updates follow equations (C11) and (C12). The extended Kalman filter uses only one point (the conditional mean) from the prior filtering density for the prediction and filtering updates.

In this paper, instead of linear-approximating the measurement equation, we directly approximate the distribution of the state vector using a set of deterministically chosen sigma points. The method is referred to as the unscented Kalman filter. See Wan and van der Merwe (2001), for example, for a detailed review of the unscented Kalman filter. Specifically, let p be the number of states, $\delta > 0$ be a control parameter, and Ω_i be the ith column of a matrix Ω. A set of $2p+1$ sigma vectors χ_i are generated according to the following equations:

$$\begin{aligned} \chi_{t,0} &= \widehat{V}_t, \\ \chi_{t,i} &= \widehat{V}_t \pm \sqrt{(p+\delta)(\widehat{\Sigma}_t)_j}, \quad j=1,\cdots,p; \quad i=1,\cdots,2p, \end{aligned} \tag{C16}$$

with corresponding weights w_i given by

$$w_0 = \delta/(p+\delta), \quad w_i = 1/[2(p+\delta)], \quad j=1,\cdots,2p. \tag{C17}$$

We can regard these sigma vectors as forming a discrete distribution with w_i being the corresponding probabilities. We can verify that the mean, covariance, skewness, and kurtosis of this distribution are $\widehat{X}_t$, $\widehat{V}_t + \mathcal{G}_t$, 0, and $p+\delta$ respectively. Thus, we can use the control parameter δ to accommodate conditional non-normalities in the state propagation equation.

Given the sigma points, the prediction on the state V and its covariance are given by

$$\overline{V}_{t+1} = \textstyle\sum_{i=0}^{2p} w_i(\Phi\chi_{t,i}), \quad \overline{\Sigma}_{t+1} = \textstyle\sum_{i=0}^{2p} w_i(\Phi\chi_{t,i} - \overline{V}_{t+1})(\Phi\chi_{t,i} - \overline{V}_{t+1})^\top + \mathcal{G}_t. \tag{C18}$$

We re-draw the sigma points $\overline{\chi}_{t+1,i}$ based on the predicted mean $\overline{V}_{t+1}$ and covariance $\overline{\Sigma}_{t+1}$ in equation (C18), and compute the predictions of the measurement $\overline{y}_{t+1}$, its variance $\overline{\Omega}_{t+1}$, and its covariance with the state $\overline{S}_{t+1}$ as,

$$\begin{aligned} \overline{y}_{t+1} &= \textstyle\sum_{i=0}^{2p} w_i O\left[\overline{\chi}_{t+1,i};\Theta\right], \\ \overline{\Omega}_{t+1} &= \textstyle\sum_{i=0}^{2p} w_i \left[O\left[\overline{\chi}_{t+1,i};\Theta\right] - \overline{y}_{t+1}\right]\left(O\left[\overline{\chi}_{t+1,i};\Theta\right] - \overline{y}_{t+1}\right)^\top + \sigma_e^2 I, \\ \overline{S}_{t+1} &= \textstyle\sum_{i=0}^{2p} w_i \left(\overline{\chi}_{t+1,i} - \overline{X}_{t+1}\right)\left(O\left[\overline{\chi}_{t+1,i};\Theta\right] - \overline{y}_{t+1}\right)^\top. \end{aligned} \tag{C19}$$

The filtering updates are

$$\begin{aligned} \widehat{V}_{t+1} &= \overline{V}_{t+1} + K_{t+1}\left(y_{t+1} - \overline{y}_{t+1}\right), \\ \widehat{\Sigma}_{t+1} &= \overline{\Sigma}_{t+1} - K_{t+1}\overline{\Omega}_{t+1}K_{t+1}^\top, \end{aligned} \tag{C20}$$

with the Kalman gain given by,

$$K_{t+1} = \overline{S}_{t+1}\left(\overline{\Omega}_{t+1}\right)^{-1}.$$

To estimate the model parameters, we define the log-likelihood for each day's observation assuming that the forecasting errors are normally distributed:

$$l_t^O(\Theta) = -\frac{1}{2}\log\left|\overline{\Omega}_t\right| - \frac{1}{2}\left(\left(y_t - \overline{y}_t\right)^\top \left(\overline{\Omega}_t\right)^{-1}\left(y_t - \overline{y}_t\right)\right). \tag{C21}$$

Appendix D. Return characteristic functions and the log likelihood

To be consistent with the weekly sampling of the index options, we also construct the log likelihood of weekly returns on the Nasdaq index. In discrete time notation, let $\ln S_t/S_{t-1}$ denote the log index return at week t over the previous week. To derive the likelihood function of index return at time t conditional on the filtered volatility rate $\widehat{V}_{t-1}$ from the option prices at time $(t-1)$, we rely on the tractability of the return characteristic function under the statistical measure $\mathbb{P}$:

$$\phi^{\mathbb{P}}[u] \quad \equiv \quad \mathbb{E}_{t-1}^{\mathbb{P}}\left(e^{iu(\ln S_t/S_{t-1})}\right), \qquad u \in \mathbb{R}. \tag{D22}$$

Given the instantaneous return risk premium derived in Theorem 1, we can rewrite the statistical dynamics of the weekly index return as,

$$\ln S_t/S_{t-1} \quad = \quad r\tau + \eta\mathcal{T}_{t-\tau,t} + W^s_{\mathcal{T}_{t-\tau,t}} + J^+_{\mathcal{T}_{t-\tau,t}} + J^-_{\mathcal{T}_{t-\tau,t}} - \xi\mathcal{T}_{t-\tau,t}, \tag{D23}$$

where $\tau = 7/365$ denotes the weekly interval of the index return. The characteristic function has a solution analogous to its risk-neutral counterpart in (B7),

$$\begin{aligned}\phi^{\mathbb{P}}[u] \;=\; & \exp\left(iur\tau - \frac{2\kappa\theta}{\omega^2}\ln\left(1 - \frac{\vartheta^* - \kappa^{\mathbb{N}}}{2\,\vartheta^*}\left(1 - e^{-\vartheta^*\tau}\right)\right) - \frac{\kappa\theta}{\omega^2}(\vartheta^* - \kappa^{\mathbb{N}})\tau \right. \\ & \left. - \frac{2\psi\left(1 - e^{-\vartheta^*\tau}\right)V_{t-1}}{2\vartheta^* - \left(\vartheta^* - \kappa^{\mathbb{N}}\right)\left(1 - e^{-\vartheta^*\tau}\right)}\right), \end{aligned} \tag{D24}$$

where $\vartheta^* \equiv \sqrt{(\kappa^{\mathbf{N}})^2 + 2\omega^2\psi}$, $\kappa^{\mathbb{N}} = \kappa - iu\omega\rho$, and $\psi[u] = \frac{1}{2}u^2 - k_{J^+}[iu] - k_{J-}[iu] + iu(\xi - \eta)$.

To compute the log likelihood of the weekly index return, we first replace V_{t-1} in equation (D24) with $\widehat{V}_{t-1}$, the filtered value from the option prices from the previous week $t-1$. Then, we apply fast Fourier inversion to the characteristic function to generate the return density numerically. Taking logarithms on the numerical density results in the weekly log likelihood that we use for the maximum likelihood estimation. □

References

Abreu, D., Brunnermeier, M., 2003. Bubbles and crashes. Econometrica 71, 173–204.

Allen, F., Gorton, G., 1993. Churning bubbles. Review of Economic Studies 60, 813–836.

Andersen, T., Benzoni, L., Lund, J., 2002. An empirical investigation of continuous-time equity return models. Journal of Finance 57, 1239–1284.

Bakshi, G., Carr, P., Wu, L., 2008. Stochastic risk premiums, stochastic skewness in currency options, and stochastic discount factors in international economies. Journal of Financial Economics 87, 132–156.

Bakshi, G., Kapadia, N., 2003. Delta-hedged gains and the negative market volatility risk premium. Review of Financial Studies 16, 527–566.

Battalio, R., Schultz, P., 2006. Options and the bubble. Journal of Finance 61, 2071–2102.

Bekaert, G., Wu, G., 2000. Asymmetric volatilities and risk in equity markets. Review of Financial Studies 13, 1–42.

Black, F., 1976. Studies of stock price volatility changes. In: Proceedings of the 1976 American Statistical Association, Business and Economical Statistics Section American Statistical Association, Alexandria, VA.

Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics 31, 307–327.

Bondarenko, O., 2004. Market price of variance risk and performance of hedge funds. Working paper. University of Illinois at Chicago.

Brunnermeier, M., Nagel, S., 2004. Hedge funds and the technology bubble. Journal of Finance 59, 2013–2040.

Campbell, J. Y., Hentschel, L., 1992. No news is good news: An asymmetric model of changing volatility in stock returns. Review of Economic Studies 31, 281–318.

Campbell, J. Y., Kyle, A. S., 1993. Smart money, noise trading and stock price behavior. Review of Economic Studies 60, 1–34.

Carr, P., Geman, H., Madan, D., Yor, M., 2002. The fine structure of asset returns: An empirical investigation. Journal of Business 75, 305–332.

Carr, P., Madan, D., 1999. Option valuation using the fast Fourier transform. Journal of Computational Finance 2, 61–73.

Carr, P., Wu, L., 2004. Time-changed Lévy processes and option pricing. Journal of Financial Economics 71, 113–141.

Carr, P., Wu, L., 2008. Variance risk premiums. Review of Financial Studies forthcoming.

Cox, J., Huang, C., 1989. Optimal consumption and portfolio policies when asset prices follow a diffusion process. Journal of Economic Theory 49, 33–83.

Dhar, R., Goetzmann, W. N., 2005. Bubble investors: What were they thinking?. Yale ICF working paper No. 05-01. Yale School of Management.

Engle, R., 1982. Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation. Econometrica 50, 987–1008.

French, K. R., with G. William Schwert, Stambaugh, R., 1987. Expected stock returns and volatility. Journal of Financial Economics 19, 3–29.

Garber, P., 1979. Tulipmania. Journal of Political Economy 97, 535–560.

Gilles, C., Leroy, S., 1992. Bubbles and charges. International Economic Review 33, 323–339.

Griffin, J., Harris, J., Topaloglu, S., 2004. Who drove and burst the tech bubble?. Working paper. University of Texas.

Harrison, M., Kreps, D., 1978. Speculative investor behavior in a stock market with heteogeneous expectations. Quarterly Journal of Economics 92, 323–336.

Harrison, M., Kreps, D., 1979. Martingales and arbitrage in multiperiod securities markets. Journal of Economic Theory 20, 381–408.

Haugen, R. A., Talmor, E., Torous, W. N., 1991. The effect of volatility changes on the level of stock prices and subsequent expected returns. Journal of Finance 46, 985–1007.

Hong, H., Scheinkman, J., Xiong, W., 2006. Asset float and speculative bubbles. Journal of Finance 61, 1073–1117.

Kindleberger, C., 1978. Manias, Panics, and Crashes: A History of Financial Crises. Wiley, New York.

Kogan, L., Ross, S., Wang, J., Westerfield, M., 2006. The price impact and survival of irrational traders. Journal of Finance 61, 195–229.

Küchler, U., Sørensen, M., 1997. Exponential Families of Stochastic Processes. Springer,, New York.

Lakonishok, J., Lee, I., Pearson, N., Poteshman, A., 2006. Option market activity. Review of Financial Studies (forthcoming).

Leland, H. E., 1980. Who should buy portfolio insurance?. Journal of Finance 35, 581–594.

Leroy, S., 2004. Rational exuberance. Journal of Economic Literature 52, 783–804.

Merton, R. C., 1976. Option pricing when underlying stock returns are discontinuous. Journal of Financial Economics 3, 125–144.

Ofek, E., Richardson, M., 2002. The valuation and market rationality of internet stock prices. Oxford Review of Economic Policy 18, 265–287.

Ofek, E., Richardson, M., 2003. Dotcom mania: The rise and fall of internet stock prices. Journal of Finance 58, 1113–1137.

Pástor, Ľ., Veronesi, P., 2005. Technological revolutions and stock prices. NBER working paper 11876 University of Chicago.

Pástor, Ľ., Veronesi, P., 2006. Was there a Nasdaq bubble in the late 1990s?. Journal of Financial Economics 81, 61–100.

Santos, M., Woodford, M., 1997. Rational asset pricing bubbles. Econometrica 65, 19–57.

Sato, K.-I., 1999. Lévy Processes and Infinitely Divisible Distributions. Cambridge University Press,, Cambridge.

Scheinkman, J., Xiong, W., 2003. Overconfidence and speculative bubbles. Journal of Political Economy 111, 1183–1219.

Schwert, G. W., 2002. Stock volatility in the new millennium: How wacky is Nasdaq?. Journal of Monetary Economics 49, 3–26.

Shefrin, H., 2005. A Behavioral Approach to Asset Pricing. Academic Press,, New York.

Shiller, R., 2000. Irrational Exuberance. Princeton University Press,, Princeton, New Jersey.

Shleifer, A., 2000. Inefficient Markets: An Introduction to Behavioral Finance. Oxford University Press,, Oxford, UK.

Temin, P., Voth, H., 2004. Riding the south sea bubble. American Economic Review 94, 1654–1668.

Wan, E. A., van der Merwe, R., 2001. The unscented Kalman filter. In: Haykin, S. (Eds.), Kalman Filtering and Neural Networks. Wiley & Sons Publishing, New York.

Wu, L., 2006. Dampened power law: Reconciling the tail behavior of financial asset returns. Journal of Business 79, 1445–1474.

Table 1

Maximum likelihood estimates of model parameters under constant market prices of risk

Entries report the maximum likelihood estimates of the model parameters under the assumption of constant market prices of risks. The absolute magnitudes of t-statistics are shown in parentheses. The estimation uses weekly data on Nasdaq 100 tracking stock returns and option prices from March 17, 1999 to February 19, 2003. In our approach, model parameters are chosen to maximize the joint conditional log likelihoods from both options and index returns,

$$\max_{\Theta} \mathcal{L}[\Theta] \equiv \sum_{t=1}^{T} \left(l_t^O[\Theta] + l_t^s[\Theta] \right),$$

where we assume conditional independence between the options forecasting errors and returns. We build the log likelihood from options under the assumption of normally distributed forecasting errors,

$$l_t^O[\Theta] = -\frac{1}{2} \ln \left|\overline{\Omega}_t\right| - \frac{1}{2} \left((y_t - \overline{y}_t)^\top \left(\overline{\Omega}_t\right)^{-1} (y_t - \overline{y}_t) \right),$$

and l_t^s is constructed based on the return characteristic function displayed in equation (D24). An extended version of the Kalman filter is adopted to handle the nonlinearity in the measurement equation.

Return Risk Dynamics, V_t:			**Jump Risk Structure, $\pi[x]$:**			**Market Prices of Risks:**		
κ	6.2608	(14.82)	λ	1.4198	(34.35)	γ_v	-3.0556	(3.84)
θ	0.1949	(21.56)	β_+	25.4845	(0.39)	$\beta_+^{\mathbb{Q}}$	4.3234	(162.83)
ω	0.5603	(34.37)	β_-	16.5948	(0.35)	$\beta_-^{\mathbb{Q}}$	1.8810	(48.56)
ρ	-0.8878	(41.73)						

Table 2

Maximum likelihood estimates of model parameters under time-varying market price of risks

Entries report the maximum likelihood estimates of the model parameters and the absolute magnitudes of t-statistics (in parentheses) with time-varying market prices of risks. Let $z_t \equiv (\gamma_v, \ln \beta_+^{\mathbb{Q}}, \ln \beta_-^{\mathbb{Q}})$ to denote the vector with dynamics:

$$z_t = \theta_z(1-\phi) + \phi z_{t-1} + \sigma_x x_t,$$

where θ_z denotes the long-run mean, the scaler ϕ captures the mean-reverting tendency of z_t, and x_t embodies time-t information relevant for updating market prices of risks. We normalize x_t to have zero mean and identity covariance matrix, with σ_x being a constant scalar. Our estimation uses weekly data on Nasdaq 100 tracking stock returns and option prices from March 17, 1999 to February 19, 2003. In our approach, model parameters are chosen to maximize the joint conditional log likelihoods from both options and index returns,

$$\max_{\Theta} \mathcal{L}[\Theta] \equiv \sum_{t=1}^{T} \left(l_t^O[\Theta] + l_t^s[\Theta] \right),$$

where we assume conditional independence between the options forecasting errors and returns. We build the log likelihood from options under the assumption of normally distributed forecasting errors,

$$l_t^O[\Theta] = -\frac{1}{2}\ln\left|\overline{\Omega}_t\right| - \frac{1}{2}\left((y_t - \overline{y}_t)^\top \left(\overline{\Omega}_t\right)^{-1}(y_t - \overline{y}_t)\right),$$

and l_t^s is constructed based on the return characteristic function displayed in equation (D24). An extended version of the Kalman filter is adopted to handle the nonlinearity in the measurement equation. The three parameters $(\beta_+^{\mathbb{Q}}, \beta_-^{\mathbb{Q}}, \gamma_v)$ reported in the third panel (labeled "Market Price of Risk Dynamics") represent their respective long-run means (θ_z).

Return Risk Dynamics, V_t			**Jump Risk Structure, $\pi[x]$:**			**Market Price of Risk Dynamics, z_t:**		
κ	7.0381	(20.14)	λ	1.5135	(53.35)	γ_v	-1.9626	(3.51)
θ	0.2052	(24.17)	β_+	33.2302	(0.31)	$\beta_+^{\mathbb{Q}}$	5.1172	(12.16)
ω	0.6429	(54.68)	β_-	21.2258	(0.35)	$\beta_-^{\mathbb{Q}}$	1.7022	(5.47)
ρ	-0.8489	(72.94)				ϕ	0.9136	(158.04)
						σ_x^2	0.0939	(26.05)

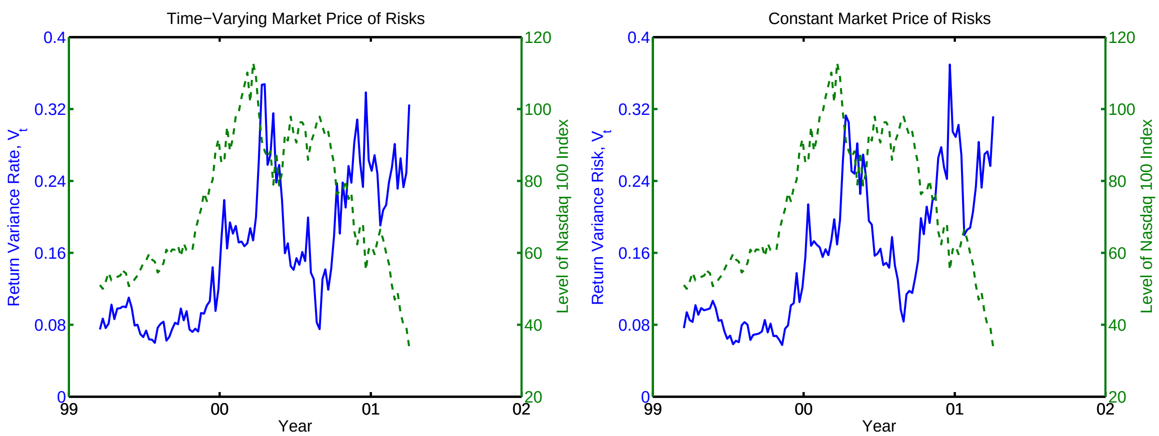

Figure 1. Return Volatility Risk versus Nasdaq Valuations

This graph plots the weekly time series of the return volatility rate V_t (solid lines) and the level of the Nasdaq 100 index tracking stock (dashed lines). The scale for the volatility rate is on the left side and the scale for the index is on the right side. The volatility rate time series in the left panel are extracted from the estimation with constant market prices of risks and that in the right panel are extracted from the estimation with time-varying market prices (see Section 3 and Appendix C).

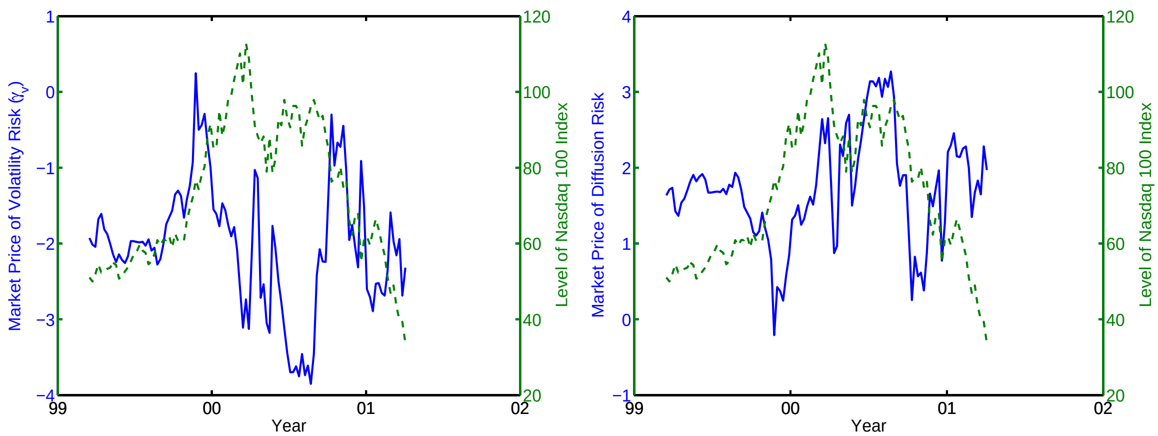

Figure 2. Time-Varying Market Price of Volatility Risk and Diffusion Risk Premium

The solid line in the left panel plots the weekly time series of the market price of the diffusive volatility risk, γ_v. The solid line in the right panel plots the time series of return risk premium per unit risk attributable to the diffusive risk η_d. The dashed lines in both panels plot the Nasdaq 100 index tracking stock closing price. The procedure for extracting γ_v is based on the estimation approach described in Section 3 and η_d is determined accordingly as $\gamma_s + \gamma_v \rho$.

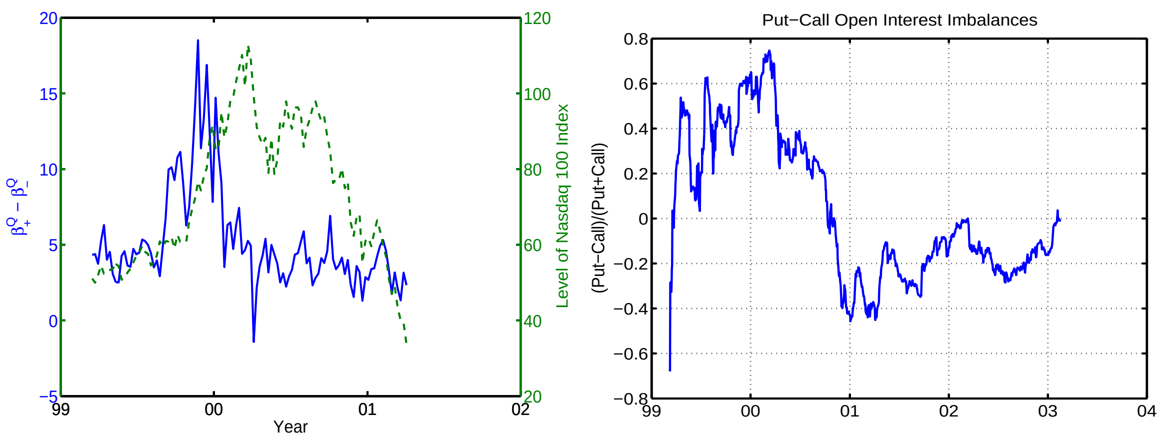

Figure 3. Market Price of Jump Risk and Open Interest

The solid line in the left panel plots the weekly time series of $\beta_+^{\mathbb{Q}} - \beta_-^{\mathbb{Q}}$, which measures the value difference between out-of-the-money put and call options. The solid line in the right panel plots the time series of the relative imbalance in the open interest between all put options and all call options on the Nasdaq 100 index tracking stock. The imbalance is defined as the difference between put and call options over the sum of calls and puts. In both panels, the dashed line is the closing level of the Nasdaq 100 index tracking stock.