Journal of Monetary Economics 60 (2013) 203–220

Contents lists available at SciVerse ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

Macroeconomic determinants of stock volatility and volatility premiums

Valentina Corradi [a], Walter Distaso [b], Antonio Mele [c,*]

[a] University of Warwick, United Kingdom
[b] Imperial College Business School, United Kingdom
[c] University of Lugano, Swiss Finance Institute, Via Buffi 13, 6900 Lugano, Switzerland

ARTICLE INFO

Article history:
Received 9 October 2010
Received in revised form
20 October 2012
Accepted 22 October 2012
Available online 1 November 2012

ABSTRACT

How does stock market volatility relate to the business cycle? We develop, and estimate, a no-arbitrage model, and find that (i) the level and fluctuations of stock volatility are largely explained by business cycle factors and (ii) some unobserved factor contributes to nearly 20% to the overall variation in volatility, although not to its ups and downs. Instead, this "volatility of volatility" relates to the business cycle. Finally, volatility risk-premiums are strongly countercyclical, even more than stock volatility, and partially explain the large swings of the VIX index during the 2007–2009 subprime crisis, which our model captures in out-of-sample experiments.

1. Introduction

Understanding the origins of stock market volatility has long been a topic of considerable interest to both policy makers and market practitioners. Policy makers are interested in the main determinants of volatility and in its spillover effects on real activity. Market practitioners are interested in the effects volatility exerts on the pricing and hedging of plain vanilla options and more exotic derivatives. In both cases, forecasting stock market volatility constitutes a formidable challenge but also a fundamental instrument to manage the risks faced by these institutions.

Many available models use latent factors to explain the dynamics of stock market volatility. For example, in the celebrated Heston's (1993) model, stock volatility is exogenously driven by some unobservable factor correlated with the asset returns. Yet such an unobservable factor does not bear an economic interpretation. Moreover, the model implies, by assumption, that volatility cannot be forecasted by macroeconomic factors such as industrial production or inflation. This circumstance is counterfactual. Indeed, there is strong evidence that stock market volatility has a very pronounced business cycle pattern, being higher during recessions than during expansions; see, e.g., Schwert (1989a,b), Hamilton and Lin (1996), or Brandt and Kang (2004).

In this paper, we develop a no-arbitrage model where stock market volatility is explicitly related to a number of macroeconomic and unobservable factors. The distinctive feature of this model is that stock volatility is linked to these factors by no-arbitrage restrictions. The model is also analytically convenient: under fairly standard conditions on the dynamics of the factors and risk-aversion corrections, our model is solved in closed-form, and is amenable to empirical work.

We use the model to quantitatively assess how market volatility and volatility-related risk-premiums change in response to business cycle conditions. Our model fully captures the procyclical nature of aggregate returns and the countercyclical behavior

* Corresponding author. Tel.: +41 586664498; fax: +41 586664734
E-mail address: antonio.mele@usi.ch (A. Mele).

0304-3932/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jmoneco.2012.10.019

of stock volatility that we have been seeing in the data for a long time. It makes a fundamental prediction: macroeconomic factors can explain nearly 75% of the variation in the overall stock volatility. At the same time, our model, rigorously estimated through simulation-based inference methods, shows that the presence of some unobservable and persistent factor is needed to sustain the level of stock volatility that matches its empirical counterpart. Moreover, our model reveals that macroeconomic factors substantially help explain the variability of stock volatility around its level—the volatility of volatility. That such a "vol-vol" might be related to the business cycle is indeed a plausible hypothesis, although clearly, the ups and downs stock volatility experiences over the business cycle are a prediction of the model in line with the data, not a restriction imposed while estimating the model. Such a new property we uncover, and model, brings practical implications. For example, business cycle forecasters might learn that not only does stock market volatility have predicting power, as discussed below; "vol-vol" is also a potential predictor of the business cycle.

Our second set of results relates to volatility-related risk-premiums. The volatility risk-premium is the difference between the expectation of future market volatility under the risk-neutral and the true probability. It quantifies how a representative agent is willing to pay to be ensured against the event that volatility will raise beyond his own expectations. Thus, it is a very intuitive and general measure of risk-aversion. We find that this volatility risk-premium is strongly countercyclical, even more so than stock volatility. Precisely, volatility risk-premiums are typically not very volatile, although in bad times, they may increase to extremely high levels, and quite quickly. We undertake a stress test of the model over a particularly uncertain period, which includes the 2007–2009 subprime turmoil. Ours is a stress test, as (i) we estimate the model using post-war data up to 2006 and (ii) feed the previously estimated model with macroeconomic data related to the subprime crisis. We compare the model's predictions for the crisis with the actual behavior of both stock volatility and the new VIX index, maintained by the Chicago Board Options Exchange (CBOE), which is, theoretically, the risk-adjusted expectation of future volatility within one month. The model tracks the dramatic movements in this index, and predicts that countercyclical volatility risk-premiums are largely responsible for the large swings in the VIX occurred during the crisis. In fact, we show that over this crisis, as well as in previous recessions, movements in the VIX index are determined by changes in such countercyclical risk-premiums, not by changes in the expected volatility.

1.1. Related literature

1.1.1. Stock volatility and volatility risk-premiums

The cyclical properties of aggregate stock market volatility have been the focus of the recent empirical research, although early work relating stock volatility to macroeconomic variables dates back to King et al. (1994), who rely on a no-arbitrage model. In a comprehensive international study, Engle and Rangel (2008) find that high frequency aggregate stock volatility has both a short-run and long-run component, and suggest that the long-run component is related to the business cycle. Adrian and Rosenberg (2008) show that the short- and long-run components of aggregate volatility are both priced, cross-sectionally. They also relate the long-run component of aggregate volatility to the business cycle. Finally, Campbell et al. (2001), Bloom (2009), Bloom et al. (2009) and Fornari and Mele (2010) show that capital market uncertainty helps explain future fluctuations in real economic activity. Our focus on volatility risk-premiums relates, instead, to the seminal work of Dumas (1995), Bakshi and Madan (2000), Britten-Jones and Neuberger (2000), and Carr and Madan (2001), which has more recently stimulated an increasing interest in these premiums dynamics and determinants (see, for example, Bakshi and Madan, 2006; Carr and Wu, 2009). Notably, in seminal work, Bollerslev and Zhou (2006) and Bollerslev et al. (2011) unveil a strong relation between volatility risk-premiums and a number of macroeconomic factors.

Our contribution hinges upon, and expands, over this growing literature, in that we formulate and estimate a fully-specified no-arbitrage model relating the dynamics of stock volatility and volatility risk-premiums to business cycle, and additional unobservable, factors. With the exception of King et al. (1994) and Adrian and Rosenberg (2008), who still have a focus different from ours, the predicting relations in the previous papers, while certainly useful, are still part of reduced-form statistical models. In out-of-sample experiments of the subprime crisis, our no-arbitrage framework will be shown to be considerably richer than that based on predictive linear regressions. For example, compared to our model's predictions about stock volatility and the VIX index, predictions from linear regressions are substantially flat over the subprime crisis.

The only antecedent to our paper is Bollerslev et al. (2009), who develop a consumption-based rationale for volatility risk-premiums, although then, the authors use this rationale only as a guidance to the estimation of reduced-form predictability regressions conditioned on the volatility risk-premium. In recent independent work discussed below, Drechsler and Yaron (2011) investigate the properties of the volatility risk-premium, implied by a calibrated consumption-based model with long-run risks. The authors, however, are not concerned with the cross-equation restrictions relating the volatility risk-premium to state variables driving low frequency stock market fluctuations which, instead, constitute the central topic of our paper.

1.1.2. No-arbitrage regressions

In recent years, there has been a significant surge of interest in consumption-based explanations of aggregate stock market volatility (see, for example, Campbell and Cochrane, 1999; Bansal and Yaron, 2004, Tauchen, 2005, Mele, 2007, or the two surveys in Campbell, 2003 and Mehra and Prescott, 2003). These explanations are important because they highlight the main economic mechanisms through which markets and preferences affect equilibrium asset prices and, hence, stock volatility. In our framework, cross-equations restrictions arise through the weaker requirement of absence of

arbitrage opportunities. In this respect, our approach is similar in spirit to the "no-arbitrage" vector autoregressions introduced in the term-structure literature by Ang and Piazzesi (2003) and Ang et al. (2006). Similarly as in those papers, we specify an analytically convenient pricing kernel affected by some macroeconomic factors, although we do not directly relate these to, say, markets, preferences or technology.

Our model works quite simply. We exogenously specify the joint dynamics of a number of macroeconomic and unobservable factors. We assume that the asset payoffs and the risk-premiums required by agents to be compensated for the fluctuations of the factors, are essentially affine functions of these factors, along the lines of Duffee (2002). We show that the resulting no-arbitrage stock price is affine in the factors. Our model does not allow for jumps or other market micro-structure effects, as our main focus is to model low frequency movements in the aggregate stock volatility and volatility risk-premiums, through the use of macroeconomic and unobservable factors. Our estimation results, obtained through data sampled at monthly frequency, are unlikely to be affected by measurement noise or jumps, say. In the related work, Carr and Wu (2009), Todorov (2010), Drechsler and Yaron (2011), and Todorov and Tauchen (2011) do allow for the presence of jumps, although they do not analyze the relations between macroeconomic variables and aggregate volatility or volatility risk-premiums, which we do here.

1.2. Estimation strategy, and plan of the paper

In standard stochastic volatility models such as that in Heston (1993), volatility is driven by factors, which are not necessarily the same as those affecting the stock price—volatility is exogenous in these models. In our no-arbitrage model, volatility is endogenous, relating to a number of risks affecting (i) macroeconomic developments, (ii) unobserved factors and (iii) the very same asset returns—these risks affect both asset returns and volatility. To identify the premium required to bear the risk of volatility, we exploit derivative data, related to the new VIX index.

We implement a three-step estimation procedure that relies on simulation-based inference methods. In the first step, we estimate the parameters underlying the macroeconomic factors. In the second step, we use data on a broad stock market index, and the macroeconomic factors, and estimate reduced-form parameters linking the stock market index to the macroeconomic factors and the third unobservable factor, as well as the parameters underlying the dynamics of the unobservable factor. In the third step, we use data on the new VIX index, and the macroeconomic factors, and estimate the risk-premiums parameters. We implement these steps by matching model-based moments and impulse response functions to their empirical counterparts, relating to macroeconomic factors, realized returns, realized volatility and the VIX index. We develop, and utilize, a theory to consistently estimate the standard errors through block-bootstrap methods.

The remainder of the paper is organized as follows. In Section 2 we develop a no-arbitrage model for the stock price, stock volatility and volatility-related risk-premiums. Section 3 illustrates the estimation strategy. Section 4 presents our empirical results. Section 5 concludes, and the Supplementary material contains an Appendix with technical details omitted from the main text.

2. The model

We develop a model where aggregate stock returns and volatility are tied up to macroeconomic developments and one unobservable factor. It is a three-factor model solved in closed form, a special case of a general multifactor model in Appendix A of the Supplementary material.[1]

2.1. The macroeconomic environment

We consider a model with one unobservable factor and two additional factors affecting the development of two aggregate macroeconomic variables, inflation and industrial production growth, and the stock market. Let $\boldsymbol{y}(t)=(y_1(t)\ y_2(t)\ y_3(t))$ be a vector-valued process, where $y_1(t)$ and $y_2(t)$ denote two observable factors, defined as $\ln(\mathrm{CPI}_t/\mathrm{CPI}_{t-12})=\ln y_1(t)$ and $\ln(\mathrm{IP}_t/\mathrm{IP}_{t-12})=\ln y_2(t)$, where CPI_t and IP_t are the consumer price index and industrial production as of month t, as further explained in Section 4.1. In Section 4.1, we also discuss the role these two macroeconomic factors have played in asset pricing. We also assume that a third, and unobservable, factor, $y_3(t)$, affects the stock price, but not the two macroeconomic aggregates, CPI_t and IP_t. Finally, we assume that the two macroeconomic factors do not affect the unobservable factor y_3, although we allow for simultaneous feedback effects between inflation and industrial production growth, as explained below. The factors y_j are solution to

$$\mathrm{d}y_j(t)=[\kappa_j(\mu_j-y_j(t))+\overline{\kappa}_j(\overline{\mu}_j-\overline{y}_j(t))]\,\mathrm{d}t+\sqrt{\alpha_j+\beta_j y_j(t)}\,\mathrm{d}W_j(t),\quad j=1,2,3, \tag{1}$$

where $W_j(t)$ are standard Brownian motions, $\overline{\mu}_1\equiv\mu_2, \overline{y}_1(t)\equiv y_2(t), \overline{\mu}_2\equiv\mu_1, \overline{y}_2(t)=y_1(t), \overline{\kappa}_3\equiv\overline{\mu}_3\equiv\overline{y}_3(t)\equiv 0$ and, finally, Greek letters denote constant parameters. The two parameters, κ_1 and κ_2, are the speed of adjustment of inflation and industrial production growth towards their long run means, μ_1 and μ_2, and $\overline{\kappa}_1$ and $\overline{\kappa}_2$ are the feedback parameters. Appendix A of the

[1] Supplemental material to this paper can be found on the journal website, or on the corresponding author's website, www.antoniomele.org.

Supplementary material reviews conditions guaranteeing Eq. (1) is well-defined, which we use as constraints whilst estimating the model.

We assume that asset prices, (i) respond to movements in the factors affecting macroeconomic conditions and (ii) reflect a long-run trend in the asset payoffs. Precisely, we model the instantaneous dividends paid off by the asset at time t, $\mathrm{Div}(t)$ say, as the product of a stochastic trend, times a stationary component, as follows:

$$\mathrm{Div}(t) = G(t)\delta(\boldsymbol{y}(t)), \tag{2}$$

where $\delta(\boldsymbol{y})$ satisfies, for four constants δ_0 and $(\delta_j)_{j=1}^3$

$$\delta(\boldsymbol{y}) = \delta_0 + \delta_1 y_1 + \delta_2 y_2 + \delta_3 y_3, \tag{3}$$

and $G(t)$ is a geometric Brownian motion with drift g and volatility σ_G,

$$\frac{\mathrm{d}G(t)}{G(t)} = g\,\mathrm{d}t + \sigma_G\,\mathrm{d}W_G(t), \quad G(0) \equiv 1, \tag{4}$$

and $W_G(t)$ is a Brownian motion uncorrelated with the Brownian motions in Eq. (1). The rationale behind the assumption in Eq. (2) is to disentangle secular, yet stochastic, dividend growth, captured by $G(t)$, from short-run fluctuations of the dividend process, arising from business cycles, and captured by $\delta(\boldsymbol{y}(t))$. This assumption implies that the asset price displays a similar property, being driven by a secular, growth component, and an additional, short-run component related to macroeconomic developments, as we now explain.

2.2. No-arbitrage

We model the pricing kernel, or the Arrow–Debreu price density, in the economy. Let $\mathbb{F}(T)$ be the sigma-algebra generated by the Brownian motion $[\boldsymbol{W}(t)^\top\ W_G(t)]^\top$, $t \le T$, where $\boldsymbol{W}(t) = (W_1(t)\ W_2(t)\ W_3(t))$, and let P the associated physical probability. The Radon–Nikodym derivative of the risk-neutral probability Q with respect to P on $\mathbb{F}(T)$ is

$$\xi(T) \equiv \frac{\mathrm{d}Q}{\mathrm{d}P} = \exp\left(-\int_0^T \boldsymbol{\Lambda}(t)^\top \mathrm{d}\boldsymbol{W}(t) - \frac{1}{2}\int_0^T \|\boldsymbol{\Lambda}(t)\|^2\,\mathrm{d}t\right) \cdot \exp\left(-\lambda_G W_G(T) - \frac{1}{2}\lambda_G^2 T\right), \tag{5}$$

for some risk-premium process $\boldsymbol{\Lambda}(t)$ and constant λ_G. The interpretation of $\boldsymbol{\Lambda}(t)$ is that of a risk-premium required to compensate for the fluctuations of the factors $\boldsymbol{y}(t)$. The constant λ_G is, instead, the unit-risk premium for the stochastic fluctuations of secular growth, $G(t)$. While we model $\boldsymbol{\Lambda}(t)$ to be time-varying, we assume λ_G to be constant for analytical convenience.

The risk-premium process is taken to satisfy an "essentially affine" specification, viz.

$$\boldsymbol{\Lambda}(\boldsymbol{y}(t)) \equiv \boldsymbol{\Lambda}(t) = \boldsymbol{V}(\boldsymbol{y}(t))\boldsymbol{\lambda}_1 + \boldsymbol{V}^-(\boldsymbol{y}(t))\boldsymbol{\lambda}_2\boldsymbol{y}(t), \tag{6}$$

where $\boldsymbol{\lambda}_1 = (\lambda_{1(1)}\ \lambda_{1(2)}\ \lambda_{1(3)})$ is a parameter vector, $\boldsymbol{\lambda}_2$ is a diagonal matrix of parameters with diagonal elements equal to $\lambda_{2(j)}$, $j = 1,2,3$, $\boldsymbol{V}(\boldsymbol{y})$ is a diagonal matrix with $\sqrt{\alpha_j + \beta_j y_j}$ on its diagonal, and $\boldsymbol{V}^-(\boldsymbol{y}) : \boldsymbol{V}^-(\boldsymbol{y})\boldsymbol{V}(\boldsymbol{y}) = \boldsymbol{I}_{3\times 3}$, for all $\boldsymbol{y}$, which it does under regularity conditions spelled out in Appendix A of the Supplementary material.

The functional form for $\boldsymbol{\Lambda}$ echoes that suggested by Duffee (2002) in the term-structure literature. If $\boldsymbol{\lambda}_2 = \boldsymbol{0}_{3\times 3}$, $\boldsymbol{\Lambda}$ collapses to the "completely affine" specification introduced by Duffie and Kan (1996), where the risk-premiums in $\boldsymbol{\Lambda}$ are tied up to the volatility of the fundamentals, $\boldsymbol{V}(\boldsymbol{y})$. While it is reasonable to assume that risk-premiums link to the *volatility* of fundamentals, the specification in Eq. (6) also allows risk-premiums to relate to the *level* of the fundamentals, through the additional term $\boldsymbol{\lambda}_2\boldsymbol{y}$. Including this term is, indeed, critical to our empirical results. Consider the *total* risk-premiums process, defined as

$$\boldsymbol{\pi}(\boldsymbol{y}) = \begin{pmatrix} \pi_1(y_1) \\ \pi_2(y_2) \\ \pi_3(y_3) \end{pmatrix} \equiv \boldsymbol{V}(\boldsymbol{y})\boldsymbol{\Lambda}(\boldsymbol{y}) = \begin{pmatrix} \alpha_1\lambda_{1(1)} + (\beta_1\lambda_{1(1)} + \lambda_{2(1)})y_1 \\ \alpha_2\lambda_{1(2)} + (\beta_2\lambda_{1(2)} + \lambda_{2(2)})y_2 \\ \alpha_3\lambda_{1(3)} + (\beta_3\lambda_{1(3)} + \lambda_{2(3)})y_3 \end{pmatrix}. \tag{7}$$

Each component of $\boldsymbol{\pi}(\boldsymbol{y})$, $\pi_j(y_j)$, depends on factor y_j due to the volatility of this factor (i.e. through β_j) and, also, due to the additional parameter $\lambda_{2(j)}$. Without $\lambda_{2(j)}$, the level of the risk-premiums could not be modeled separately from their sensitivities to changes in y_j—a sensible issue we have experienced whilst estimating the model. Consider, for example, the total risk premium for growth, $\pi_2(y_2)$. The coefficient $\lambda_{1(2)}$ affects both the intercept and the slope of π_2. The inclusion of $\lambda_{2(2)}$ allows to achieve flexibility in modeling the level of $\pi_2(y_2)$ and its sensitivity with respect to changes in y_2.

Our final assumption is that the safe asset is elastically supplied such that the short-term rate r (say) is constant. Whilst real rates are not as volatile as stock returns in the data, many existing models might likely predict rates to be too volatile. For example, models with habit formation predict that the short-term rate is a function of the state, primarily due to intertemporal substitution effects. Campbell and Cochrane (1999) mitigate this issue with a well-known trick—they impose that intertemporal substitution effects are exactly offset by precautionary savings, thereby making the short-term rate constant. Additional models that cope with this challenge include those relying on non-expected utility, as in Bansal and Yaron (2004), or those with heterogeneous agents, as in Guvenen (2009), to cite a few. In this paper, r is kept constant for the purpose of keeping stock volatility tractable, as this facilitates the actual estimation of the model. How important is

this assumption, quantitatively? Mele (2007) finds that in realistically calibrated models of habit formation, large countercyclical swings of stock volatility mainly arise due to risk-premiums effects, rather than interest rate volatility. It is an open question, however, whether such a result would still hold in the economy we consider in the current paper.

2.3. Stock price and volatility

We are ready to determine the no-arbitrage stock price. As it turns out, the previous assumption on the pricing kernel and the assumption that $\delta(\cdot)$ in Eq. (3) is affine in $\boldsymbol{y}$ implies that the stock price is also affine in $\boldsymbol{y}$. Precisely, we have

$$S(G,\boldsymbol{y}) = G \cdot \left(s_0 + \sum_{j=1}^{3} s_j y_j \right), \tag{8}$$

where

$$s_0 = \frac{1}{r-g+\sigma_G \lambda_G} \left[\delta_0 + \sum_{j=1}^{3} s_j(\kappa_j \mu_j + \overline{\kappa}_j \overline{\mu}_j - \alpha_j \lambda_{1(j)}) \right], \tag{9}$$

$$s_j = \frac{\delta_j(r-g+\sigma_G\lambda_G+\kappa_i+\lambda_{1(i)}\beta_i+\lambda_{2(i)})-\delta_i\overline{\kappa}_i}{\prod_{h=1}^{2}(r-g+\sigma_G\lambda_G+\kappa_h+\lambda_{1(h)}\beta_h+\lambda_{2(h)})-\overline{\kappa}_1\overline{\kappa}_2} \quad \text{for } j,i \in \{1,2\} \text{ and } i \neq j, \tag{10}$$

$$s_3 = \frac{\delta_3}{r-g+\sigma_G\lambda_G+\kappa_3+\lambda_{1(3)}\beta_3+\lambda_{2(3)}}. \tag{11}$$

In the standard stochastic volatility literature, the asset price and, hence, its volatility, is taken as given, and volatility and volatility risk-premiums are modeled separately, as for example in the celebrated Heston's (1993) model, which take many empirical studies as a benchmark (e.g., Chernov and Ghysels, 2000; Corradi and Distaso, 2006; Garcia et al., 2011). Moreover, a recent focus in this literature is to relate volatility risk-premiums to the business cycle (e.g., Bollerslev et al., 2011). Yet, while the empirical results in these papers are ground breaking, Heston's model is not meant to capture, theoretically, the interplay between stochastic volatility, volatility risk-premiums and the business cycle.

Our model works differently, as it places restrictions on the asset price process directly, through our assumptions on the fundamentals of the economy, and the absence of arbitrage. For our model, it is the asset price that determines, endogenously, volatility, which by Eq. (1) and Eq. (8) is

$$\sigma(\boldsymbol{y}(t)) \equiv \sigma(t) = \sqrt{\sigma_G^2 + \frac{\sum_{j=1}^{3} s_j^2(\alpha_j+\beta_j y_j(t))}{(s_0+\sum_{j=1}^{3} s_j y_j(t))^2}}. \tag{12}$$

Note that the model predicts that stock volatility embeds information about risk-corrections that agents require to invest in the stock market. We shall make use of this observation in the empirical part of the paper. We now describe which measure of stock volatility we use to proceed with such a critical step of our analysis.

2.4. Arrow–Debreu adjusted volatility

In September 2003, the CBOE changed its volatility index VIX, to reflect recent advances in the option pricing literature. Given an asset price process $S(t)$ that is continuous in time (as that predicted by our model, in Eq. (8)), and all available information $\mathbb{F}(t)$ at time t, consider the economic value of the future integrated variance on a given interval $[t,t_0]$, IV_{t,t_0}, say, which is the sum of the future variances, weighted with the Arrow–Debreu state prices:

$$\mathbb{E}[IV_{t,t_0} \,|\, \mathbb{F}(t)] \equiv \int_t^{t_0} \mathbb{E}\left[\left(\frac{\mathrm{d}}{\mathrm{d}\tau} \mathrm{var}[\ln S(\tau) \,|\, \mathbb{F}(u)] \Big|_{\tau=u} \right) \Big| \mathbb{F}(t) \right] \mathrm{d}u, \tag{13}$$

where $\mathbb{E}$ is the expectation under Q. The new VIX index relies on the work of Dumas (1995), Bakshi and Madan (2000), Britten-Jones and Neuberger (2000), and Carr and Madan (2001), who showed that the risk-neutral expectation of the future integrated variance is a functional of put and call options written on the asset:

$$\mathbb{E}[IV_{t,t_0} \,|\, \mathbb{F}(t)] = 2e^{r(t_0-t)} \left[\int_0^{F(t)} \frac{P_t(t_0,K)}{K^2} \mathrm{d}K + \int_{F(t)}^{\infty} \frac{C_t(t_0,K)}{K^2} \mathrm{d}K \right] \equiv (t_0-t) \cdot \mathrm{VIX}_t^2, \tag{14}$$

where $F(t) = e^{r(t_0-t)}S(t)$ is the forward price, $C_t(t_0,K)$ and $P_t(t_0,K)$ are the prices as of time t of call and put options expiring at t_0 and struck at K, and VIX_t is the new VIX index. In contrast, our model, which relies on the Arrow–Debreu state prices in Eq. (5), predicts that the risk-neutral expectation of the integrated variance is

$$\mathbb{E}[IV_{t,t_0} \,|\, \boldsymbol{y}(t) = \boldsymbol{y}] = \int_t^{t_0} \mathbb{E}[\sigma^2(\boldsymbol{y}(u)) \,|\, \boldsymbol{y}(t) = \boldsymbol{y}] \, \mathrm{d}u \equiv (t_0-t) \cdot \mathrm{VIX}^2(\boldsymbol{y}), \tag{15}$$

where $\sigma^2(\boldsymbol{y}(t))$ is given in Eq. (12). We shall estimate the risk-premium parameters in Eq. (7) so as to match the VIX index predicted by the model, $\mathrm{VIX}(\boldsymbol{y}(t))$ in Eq. (15), to its empirical counterpart, VIX_t in Eq. (14). Finally, our model makes predictions about how the volatility risk-premium, $\mathrm{VRP}(\boldsymbol{y}(t))$ say, changes with the factors $\boldsymbol{y}(t)$ in Eq. (1)

$$\mathrm{VRP}(\boldsymbol{y}(t)) \equiv \sqrt{\frac{1}{t_0-t}}\left(\sqrt{\mathbb{E}[IV_{t,t_0}|\boldsymbol{y}(t)=\boldsymbol{y}]}-\sqrt{E[IV_{t,t_0}|\boldsymbol{y}(t)=\boldsymbol{y}]}\right), \tag{16}$$

where E denotes the expectation taken under P.

3. Statistical inference

We rely on a three-step procedure. In the first step, we estimate the parameters of the process underlying the dynamics of the two macroeconomic factors, $\boldsymbol{\phi}^\top = \left(\kappa_j,\mu_j,\alpha_j,\beta_j,\overline{\kappa}_j,\ j=1,2\right)$. In the second step, we estimate the parameters in Eq. (4), $\boldsymbol{\theta}_G^\top = (g,\sigma_G)$, the reduced-form parameters that link the asset price to the three factors in Eq. (8), and the parameters of the process for the unobserved factor, $\boldsymbol{\theta}^\top = (\kappa_3,\mu_3,\alpha_3,\beta_3,(s_j)_{j=0}^3)$, while imposing the identifiability condition that $\mu_3=1$, as explained below. In the third step, we estimate the risk-premiums parameters $\boldsymbol{\lambda}^\top = (\lambda_{1(1)},\lambda_{2(1)},\lambda_{1(2)},\lambda_{2(2)},\lambda_{1(3)},\lambda_{2(3)})$, relying on a simulation-based approximation of the model-implied VIX, which we match to the VIX index. At each of these steps, we do not have a closed form expression for the likelihood function, or for selected sets of moment conditions. For this reason, we need to rely on a simulation-based approach. Our estimation strategy relies on a hybrid of Indirect Inference (Gouriéroux et al., 1993) and the Simulated Generalized Method of Moments (Duffie and Singleton, 1993).[2]

3.1. Moment conditions for the macroeconomic factors

To simulate the factor dynamics in Eq. (1), we rely on a Milstein approximation scheme, with discrete interval Δ, say. We simulate H paths of length T of the two observable factors, and sample them at the same frequency as the available data, obtaining $y^{\phi}_{1,t,\Delta,h}$ and $y^{\phi}_{2,t,\Delta,h}$, where $y^{\phi}_{j,t,\Delta,h}$ is the value at time t taken by the j-th factor, at the h-th simulation performed with $\boldsymbol{\phi}$—the parameter vector relating to the process underlying the macroeconomic factors. Then, we estimate the following auxiliary models on both historical and simulated data[3]

$$\boldsymbol{y}_t = \boldsymbol{w}+\boldsymbol{A}\boldsymbol{y}_{t-1}+\boldsymbol{\epsilon}_t^y, \tag{17}$$

and

$$\boldsymbol{y}^{\phi}_{t,\Delta,h} = \boldsymbol{w}_h+\boldsymbol{A}_h\boldsymbol{y}^{\phi}_{t-1,\Delta,h}+\boldsymbol{\epsilon}^y_{t,h}, \tag{18}$$

where $\boldsymbol{y}_t = (y_{1,t}\ y_{2,t})^\top$, $\boldsymbol{w}$ and $\boldsymbol{A}$ denote a vector and a matrix of constants, $\boldsymbol{\epsilon}_t^y$ is a vector of errors with zero mean and (diagonal) variance–covariance matrix $\boldsymbol{C}$, and the notation for Eq. (18) for simulated data follows the same rationale as that in Eq. (17).

Next, let $\tilde{\boldsymbol{\varphi}}_T = (\tilde{\boldsymbol{\varphi}}_{1,T},\tilde{\boldsymbol{\varphi}}_{2,T},\overline{y}_1,\overline{y}_2,\hat{\sigma}_1,\hat{\sigma}_2)^\top$, where $\tilde{\boldsymbol{\varphi}}_{1,T}$ and $\tilde{\boldsymbol{\varphi}}_{2,T}$ denote the ordinary least squares (OLS, henceforth) estimators of the parameters in Eq. (17), and $\overline{y}_i$ and $\hat{\sigma}_i$ are the sample mean and standard deviation of the macroeconomic factors. Let $\tilde{\boldsymbol{\varphi}}_{T,\Delta,h}(\boldsymbol{\phi})$ be the simulated counterpart to $\tilde{\boldsymbol{\varphi}}_T$ at simulation h, including the OLS estimator of the parameters in Eq. (18), and the sample means and standard deviations of the macroeconomic factors. The estimator of $\boldsymbol{\phi}$ is

$$\hat{\boldsymbol{\phi}}_T \equiv \arg\min_{\boldsymbol{\phi}\in\boldsymbol{\Phi}_0}\left\|\frac{1}{H}\sum_{h=1}^{H}\tilde{\boldsymbol{\varphi}}_{T,\Delta,h}(\boldsymbol{\phi})-\tilde{\boldsymbol{\varphi}}_T\right\|^2, \tag{19}$$

where $\boldsymbol{\Phi}_0$ is some compact set. Appendix B in the Supplementary material develops the asymptotic theory relating to this estimator.

3.2. Moment conditions for realized returns and volatility

Data on macroeconomic factors and stock returns do not allow us to identify the structural parameters of the model. In particular, there are many combinations of $\boldsymbol{\delta} = (\delta_j)_{j=0}^3$ and $\boldsymbol{\lambda} = (\lambda_{1(j)},\lambda_{2(j)})_{j=1}^3$ in Eqs. (9)–(11), giving rise to the same asset price. In this second step, we estimate the parameters $\boldsymbol{\theta}_G = (g,\sigma_g)$ in Eq. (4), the reduced-form parameters, $(s_j)_{j=0}^3$ in Eqs. (9)–(11), and the parameters for the unobservable factor $(\kappa_3,\mu_3,\alpha_3,\beta_3)$. The parameters $\boldsymbol{\lambda}$ shall be estimated in a third

[2] The estimators we develop are not as efficient as Maximum Likelihood. Under some conditions, the methods put forward by Gallant and Tauchen (1996), Fermanian and Salanié (2004), Carrasco et al. (2007), Aït-Sahalia (2008), or Altissimo and Mele (2009) are asymptotic equivalent to Maximum Likelihood. In our context, they deliver asymptotically efficient estimators for the parameters in the first step. However, hinging upon these approaches in the remaining steps would make the two issues of unobservability of volatility and, especially, parameter estimation error considerably beyond the scope of this paper.

[3] The choice of lags for all the auxiliary models in the present section relies on the BIC criterion, and our additional concern to have non-overlapping regressors—with the exception of Eqs. (17), and a lag 6 in Eq. (23), which revealed to be empirically important. Appendix C.2 of the Supplementary material reports parameter estimates and R^2 for all these auxiliary models—both those referring to data and those implied by the model.

and final step, described in the next section. Note that, theoretically, it might be possible to collapse the second and third steps of our estimation procedure into a single one, where a combined use of data on dividends and volatility derivatives might help identify $\boldsymbol{\delta}$ and $\boldsymbol{\lambda}$. We do not pursue this approach because it revealed to be computationally prohibitive. Note, then, that our three-step methodology leads to identify the model's parameters through data relating to macroeconomic factors, and market data relating to stock returns and risk-neutral volatility. However, Appendix C.1 of the Supplementary material describes a calibration procedure relying on both the aggregate asset price and dividends, which leads us to find our three-step estimation methodology has quite reasonable implications for the dividends dynamics.

Even proceeding in this way, we cannot tell apart the loading on the unobservable factor, s_3, from the parameters underlying the dynamics of this factor $(\kappa_3,\mu_3,\alpha_3,\beta_3)$ as this is independent of the observable ones. We impose the normalization $\mu_3 \equiv 1$. We estimate $\boldsymbol{\theta}_G$ using the time-series of the low-frequency component of the real stock price growth, extracted through the Hodrick–Prescott filter with smoothing parameter equal to 14,400, given we are using monthly data (Hodrick and Prescott, 1997). We simulate H paths of length T of the unobservable factor $y_3(t)$, and the secular growth, $G(t)$, using a Milstein approximation with discrete interval Δ, and sample them at the same frequency as the data, obtaining for $\boldsymbol{\theta}_u = (\kappa_3,\alpha_3,\beta_3,s_3)$ and $\hat{\boldsymbol{\theta}}_{G,T} = (\hat{g}_T,\hat{\sigma}_{G,T})$, and simulation h, the series $y^{\theta_u}_{3,t,\Delta,h}$ and $G^{\hat{\theta}_{G,T}}_{t,\Delta,h}$. Likewise, let $S^{\theta}_{t,\Delta,h}(\hat{\boldsymbol{\theta}}_{G,T})$ be the simulated series of the stock price, when the parameters are fixed at $\boldsymbol{\theta} = (\boldsymbol{\theta}_u,(s_j)^3_{j=0})$ and $\hat{\boldsymbol{\theta}}_{G,T}$

$$\ln S^{\theta}_{t,\Delta,h}(\hat{\boldsymbol{\theta}}_{G,T}) = \ln G^{\hat{\theta}_{G,T}}_{t,\Delta,h} + \ln(s_0 + s_1 y_{1,t} + s_2 y_{2,t} + s_3 y^{\theta_u}_{3,t,\Delta,h}), \tag{20}$$

where $G^{\hat{\theta}_{G,T}}_{0,\Delta,h} \equiv 1$, as in Eq. (4). We fix the intercept, s_0, so as to make the model-implied average of the detrended stock price match its empirical counterpart: $s_0 = \overline{S}^d - s_1\overline{y}_1 - s_2\overline{y}_2 - s_3$, where $\overline{S}^d$ denotes the sample mean of the detrended stock price $S^d_t \equiv e^{-\hat{g}_T t} S_t$, S_t is the real stock price index observed at time t, and finally, $\overline{y}_1$ and $\overline{y}_2$ are the sample means of the two macroeconomic factors $y_{1,t}$ and $y_{2,t}$. Note that we simulate the stock price using the *observed* samples of $y_{1,t}$ and $y_{2,t}$, a feature of the estimation strategy that results in an improved efficiency, as discussed below.

Following Mele (2007) and Fornari and Mele (2010), we measure the volatility of the monthly continuously compounded price changes, as

$$\mathrm{Vol}_t = \sqrt{6\pi} \cdot \frac{1}{12} \sum_{i=1}^{12} \left| \ln\left(\frac{S_{t+1-i}}{S_{t-i}}\right) \right|. \tag{21}$$

Next, define yearly returns as $R_t = \ln(S_t/S_{t-12})$, and let $R^{\theta}_{t,\Delta,h}(\hat{\boldsymbol{\theta}}_{G,T})$ and $\mathrm{Vol}^{\theta}_{t,\Delta,h}(\hat{\boldsymbol{\theta}}_{G,T})$ be the simulated counterparts to R_t and Vol_t.[4]

Our estimator relies on two auxiliary models that capture the main statistical facts about stock returns and return volatility in our dataset. The auxiliary model for returns is

$$R_t = a^{\mathrm{R}} + b^{\mathrm{R}}_{1,12} y_{1,t-12} + b^{\mathrm{R}}_{2,12} y_{2,t-12} + \epsilon^{\mathrm{R}}_t, \tag{22}$$

and that for return volatility is

$$\mathrm{Vol}_t = a^{\mathrm{V}} + \sum_{i\in\{6,12,18,24,36,48\}} b^{\mathrm{V}}_i \mathrm{Vol}_{t-i} + \sum_{i\in\{12,24,36,48\}} b^{\mathrm{V}}_{1,i} y_{1,t-i} + \sum_{i\in\{12,24,36,48\}} b^{\mathrm{V}}_{2,i} y_{2,t-i} + \epsilon^{\mathrm{V}}_t. \tag{23}$$

Let $\tilde{\boldsymbol{\vartheta}}_T = (\tilde{\boldsymbol{\vartheta}}_{1,T},\tilde{\boldsymbol{\vartheta}}_{2,T},\overline{S},\overline{\mathrm{Vol}},\hat{\sigma}_{\mathrm{Vol}})^{\top}$, where $\tilde{\boldsymbol{\vartheta}}_{1,T}$ is the OLS estimate of the parameters in Eq. (22), $\tilde{\boldsymbol{\vartheta}}_{2,T}$ is the OLS estimate of the parameters in Eq. (23), $\overline{S}$ is the sample mean of the real stock price, and, finally, $\overline{\mathrm{Vol}}$ and $\hat{\sigma}_{\mathrm{Vol}}$ are the sample mean and standard deviation of stock return volatility. Let $\hat{\boldsymbol{\vartheta}}_{T,\Delta,h}(\boldsymbol{\theta},\hat{\boldsymbol{\theta}}_{G,T})$ be the simulated counterpart to $\tilde{\boldsymbol{\vartheta}}_T$ at simulation h, using $R^{\theta}_{t,\Delta,h}(\hat{\boldsymbol{\theta}}_{G,T})$ and $\mathrm{Vol}^{\theta}_{t,\Delta,h}(\hat{\boldsymbol{\theta}}_{G,T})$. The estimator of $\boldsymbol{\theta} = (\boldsymbol{\theta}_u,(s_j)^3_{j=0})$ is

$$\hat{\boldsymbol{\theta}}_T = \arg\min_{\boldsymbol{\theta}\in\boldsymbol{\Theta}_0} \left\| \frac{1}{H} \sum_{h=1}^{H} \hat{\boldsymbol{\vartheta}}_{T,\Delta,h}(\boldsymbol{\theta},\hat{\boldsymbol{\theta}}_{G,T}) - \tilde{\boldsymbol{\vartheta}}_T \right\|^2, \tag{24}$$

where $\boldsymbol{\Theta}_0$ is a compact set. As shown in detail in Appendix B of the Supplementary material, the structure of the asymptotic covariance matrix of this estimator differs from that of $\hat{\boldsymbol{\phi}}_T$ in Eq. (19), due to two reasons. First, stock price paths are simulated through Eq. (20), with secular growth parameters fixed at their estimates, $\hat{\boldsymbol{\theta}}_{G,T}$, leading to parameter estimation error, which is asymptotically accounted for. Second, ours is, in fact, a conditional simulated inference estimator, in that the simulations in Eq. (20) occur conditionally upon the sample realizations of the observable factors, $y_{1,t}$ and $y_{2,t}$. This feature of the method results in a correlation among the auxiliary parameter estimates obtained over all the simulations, and leads to an efficiency improvement over unconditional (simulated) inference.

[4] Our asymptotics are not affected by the property that the returns $R^{\theta}_{t,\Delta,h}(\cdot)$ are driven by both a $I(0)$ and a $I(-1)$ component, as our laws of large numbers or central limit theorems still apply—similarly as in the realized volatility literature, for instance, where $I(-1)$ terms arise due to microstructure effects (see, e.g., Barndorff-Nielsen et al., 2008; or Aït-Sahalia et al., 2011). Campbell and Cochrane (1999) model of habit formation or Bansal and Yaron (2004) long-run risks model are instances of models predicting a similar property for asset returns.

3.3. Estimation of the risk-premium parameters

We estimate the risk-premium parameters, $\boldsymbol{\lambda}$, by matching moments and impulse response functions of the model-based VIX, VIX($\boldsymbol{y}(t)$) in Eq. (15), to those of the model-free VIX index, VIX_t in Eq. (14), with t_0-t equal to one month. Since the new VIX index is available only since 1990, we use a sample of $\mathcal{T}$ observations in this step, with $\mathcal{T}<T$. Whilst VIX($\boldsymbol{y}(t)$) is not known in closed-form, it can be accurately approximated through simulations, as explained in Appendix B of the Supplementary material. Note, also, that in the actual computation of Eq. (15), we replace the unknown parameters, s_0 $(s_j,\kappa_j,\alpha_j,\beta_j)_{j=1}^3$, $(\overline{\kappa}_i,\mu_i)_{i=1}^2$ and σ_G, with their estimated counterparts computed in the previous two steps: $\hat{\boldsymbol{\theta}}_T$, $\hat{\boldsymbol{\phi}}_T$ and $\hat{\sigma}_{G,T}$. As in the previous step, we use the observed samples of the macroeconomic factors $y_{1,t}$,$y_{2,t}$, and simulate samples for the latent factor only. We rely on the following auxiliary model:

$$\text{VIX}_t = a^{\text{VIX}} + b^{\text{VIX}}\text{VIX}_{t-1} + \sum_{i\in\{36,48\}} b_{1,i}^{\text{VIX}} y_{1,t-i} + \sum_{i\in\{36,48\}} b_{2,i}^{\text{VIX}} y_{2,t-i} + \epsilon_t^{\text{VIX}}. \tag{25}$$

Define $\tilde{\boldsymbol{\psi}}_{\mathcal{T}} = (\tilde{\boldsymbol{\psi}}_{1,\mathcal{T}},\overline{\text{VIX}},\hat{\sigma}_{\text{VIX}})^\top$, where $\tilde{\boldsymbol{\psi}}_{1,\mathcal{T}}$ is the OLS estimator of the parameters in Eq. (25), and $\overline{\text{VIX}}$ and $\hat{\sigma}_{\text{VIX}}$ are the sample mean and standard deviation of the VIX index. Likewise, define $\hat{\boldsymbol{\psi}}_{\mathcal{T},\Delta,h}(\hat{\boldsymbol{\theta}}_T,\hat{\boldsymbol{\phi}}_T,\hat{\sigma}_{G,T},\boldsymbol{\lambda})$, the simulated counterpart to $\tilde{\boldsymbol{\psi}}_{\mathcal{T}}$ at simulation h, obtained through simulations of the model-implied index, $\text{VIX}_{t,\Delta,h}(\hat{\boldsymbol{\theta}}_T,\hat{\boldsymbol{\phi}}_T,\hat{\sigma}_{G,T},\boldsymbol{\lambda})$ say, where the paths of the two macroeconomic factors, $y_{1,t}$ and $y_{2,t}$, are fixed at their sample values. The estimator of $\boldsymbol{\lambda}$ is

$$\hat{\boldsymbol{\lambda}}_{\mathcal{T}} = \arg\min_{\boldsymbol{\lambda}\in\boldsymbol{\Lambda}_0}\left\|\frac{1}{H}\sum_{h=1}^{H}\hat{\boldsymbol{\psi}}_{\mathcal{T},\Delta,h}(\hat{\boldsymbol{\theta}}_T,\hat{\boldsymbol{\phi}}_T,\hat{\sigma}_{G,T},\boldsymbol{\lambda}) - \tilde{\boldsymbol{\psi}}_{\mathcal{T}}\right\|^2, \tag{26}$$

for some compact set $\boldsymbol{\Lambda}_0$. This estimator is, similarly as $\hat{\boldsymbol{\theta}}_T$ in Eq. (24), affected by parameter estimation error, arising because $\text{VIX}_{t,\Delta,h}(\hat{\boldsymbol{\theta}}_T,\hat{\boldsymbol{\phi}}_T,\hat{\sigma}_{G,T},\boldsymbol{\lambda})$, the model-implied VIX index, is simulated using parameters estimated in the previous two steps, $\hat{\boldsymbol{\phi}}_T$, $\hat{\boldsymbol{\theta}}_T$ and $\hat{\sigma}_{G,T}$. At the same time, the estimator $\hat{\boldsymbol{\lambda}}_{\mathcal{T}}$ in Eq. (24) is a conditionally simulated one, in that it relies on the observations of the macroeconomic factors $y_{1,t}$ and $y_{2,t}$, thereby resulting in efficiency gains.

3.4. Bootstrap standard errors

The limiting variance–covariance matrices for $\hat{\boldsymbol{\phi}}_T$ in Eq. (19), $\hat{\boldsymbol{\theta}}_T$ in Eq. (24), and $\hat{\boldsymbol{\lambda}}_{\mathcal{T}}$ in Eq. (26) are characterized in Appendix B.1 of the Supplementary material. They are not known in closed form, and must be estimated through the computation of several numerical derivatives. Moreover, our sample sizes are relatively small, compared to those we usually have access to in empirical finance, and in particular such is that available for the estimation of the risk premium parameters. We rely on bootstrap standard errors consistent for those implied by the asymptotic variance–covariance matrices for $\hat{\boldsymbol{\phi}}_T$, $\hat{\boldsymbol{\theta}}_T$ and $\hat{\boldsymbol{\lambda}}_{\mathcal{T}}$. Bootstrap standard errors are not only easier to compute, but also less prone to numerical errors, and likely to be more accurate than those based on asymptotic approximations, in finite samples. Finally, the auxiliary models we utilize are potentially misspecified, and they likely lead to a score that is not a martingale difference sequence. We appeal to the "block-bootstrap" to address this technical issue. Appendix B.2 of the Supplementary material develops results and algorithms that allow us to make use of this method within the simulation-based estimation procedure of this section.

4. Empirical analysis

This section contains empirical results. Section 4.1 describes and motivates our dataset. Section 4.2 provides model's estimates and predictions on aggregate returns and volatility, volatility risk-premiums, factor contributions and, finally, out-of-sample experiments. Appendix C of the Supplementary material contains additional results, and further details are omitted here for the sake of brevity.

4.1. Data

Our security data include the S&P 500 Compounded index, and the VIX index maintained by the CBOE. The VIX index is available daily, but only after January 1990. Our macroeconomic variables include the consumer price index (CPI), and the seasonally adjusted industrial production (IP) index for the US. Information related to the CPI and the IP indexes is made available to the market between the 19th and the 23th of every month. To possibly avoid overreaction to releases of information, we sample the S&P Compounded index and the VIX index every 25th of the month. We compute the real stock price as the ratio between the S&P index and the CPI. Our dataset, then, includes (i) monthly observations of the VIX index, from January 1990 to December 2006, for a total of 204 observations and (ii) monthly observations of the real stock price, the CPI and the IP indexes, from January 1950 to December 2006, for a total of 672 observations.

Our dataset also includes monthly observations of the University of Michigan Consumer Sentiment index, from January 1978 to December 2006 (for a total of 336 observations). Finally, we utilize additional data, from January 2007 to March 2009, to implement a stress test of how the previously estimated model would have performed over a particularly critical period. This out-of-sample period is critical for at least three reasons: first, the NBER determined that the US economy entered in a recession

in December 2007, which is the third NBER-dated recession since the creation of the new VIX index; second, this period includes the quite unique events leading to the subprime crisis; and third, both realized stock market volatility and the VIX index reached record highs, and possibly pose challenges to rational models of asset prices. Note that our out-of-sample experiments are not intended to forecast the market, stock market volatility, and the level of the VIX index. Rather, we feed the model estimated up to December 2006, with macroeconomic data (the CPI and IP indexes) available from January 2007, and compare the predictions of the model with the actual movements of the market, stock market volatility and the VIX index.

Many theoretical explanations and, in fact, the empirical evidence, would lead us to expect that asset prices are, indeed, related to variables tracking business cycles (see, e.g., Cochrane, 2005), such as the CPI and the IP growth. For example, in their seminal article relating stock returns to the macroeconomy, Chen et al. (1986) find that industrial production growth and inflation are among the most prominent priced factors. Theoretically, in standard theories of external habit formation, the pricing kernel volatility is driven by the surplus consumption ratio, defined as the percentage deviation of current consumption, C, from some habit level, H, i.e. $(C-H)/C$, which highly correlates with procyclical variables such as industrial production growth. Likewise, standard asset pricing models predict that compensation for inflation risk relates to variables that are highly correlated with inflation (e.g., Bakshi and Chen, 1996; Buraschi and Jiltsov, 2005). Mainly for computational reasons, we refrain from considering additional factors to model the linkages of the pricing kernel to the business cycle.

Fig. 1 depicts the two series $y_{1,t}$ (year-to-year gross inflation) and $y_{2,t}$ (year-to-year industrial production growth) along with NBER-dated recession events. Gross inflation is procyclical, although it peaked up during the 1975 and the 1980 recessions, as a result of the geopolitical driven oil crises that occurred in 1973 and 1979. Its volatility during the 1970s was large until the Monetary experiment of the early 1980s, although it dramatically dropped during the period following the experiment, usually referred to as the Great Moderation (e.g., Bernanke, 2004). At the same time, inflation is persistent: a Dickey–Fuller test rejects the null hypothesis of a unit root in $y_{1,t}$, although the rejection is at the marginal 95% level. The inclusion of inflation as a determinant of the pricing kernel displays one attractive feature. An old debate exists upon whether stocks provide a hedge against inflation (see, e.g.; Danthine and Donaldson, 1986). While our no-arbitrage model is silent about the general equilibrium forces underlying inflation-hedge properties of asset prices, its data-driven structure allows us to assess quite directly the relations between inflation and the stock price, returns, volatility and volatility risk-premiums.

Fig. 1 also shows that while the volatility of industrial production growth dropped during the Great Moderation, growth is still persistent, although less so than gross inflation: here, a Dickey–Fuller test rejects the null hypothesis of a unit root in $y_{2,t}$ at any conventional level. Finally, the properties of inflation and industrial production growth over our out-of-sample period, from January 2007 to March 2009, are discussed in Section 4.2.4.

4.2. Estimation results

This section contains our core empirical results. Sections 4.2.1 and 4.2.2 provide estimates for the two macroeconomic and the unobservable factors, secular growth, and the stock price processes, along with some basic model's implications on aggregate stock returns and volatility. Section 4.2.3 describes measures of "contributions" to stock volatility borne by each

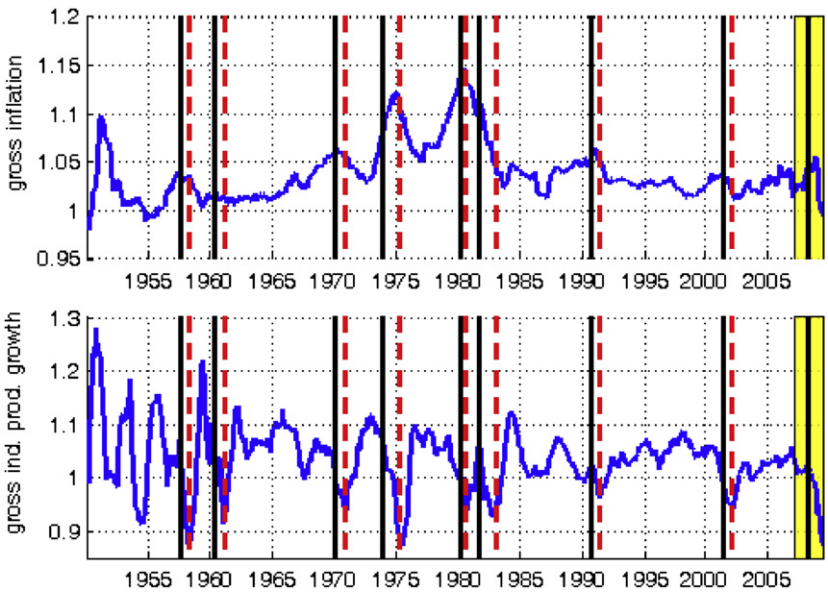

Fig. 1. This figure plots one-year, monthly gross inflation and one-year, monthly gross industrial production growth, as defined in the main text. The sample covers the period from January 1950 to December 2006. Vertical solid lines track the beginning of NBER-dated recessions and vertical dashed lines indicate the end of NBER-dated recessions. The shaded area covers the out-of-sample period, from January 2007 to March 2009.

Table 1
Parameter estimates and block-bootstrap standard errors for the joint process of the two macroeconomic factors in Eq. (1), gross inflation, $y_{1,t} \equiv \mathrm{CPI}_t/\mathrm{CPI}_{t-12} \equiv y_1(t)$ and gross industrial production growth, $y_{2,t} \equiv \mathrm{IP}_t/\mathrm{IP}_{t-12} \equiv y_2(t)$, where CPI_t is the consumer price index as of month t, IP_t is the real, seasonally adjusted industrial production index as of month t, and the parameter vector is $\boldsymbol{\phi}^\top = (\kappa_j, \mu_j, \alpha_j, \beta_j, \overline{\kappa}_j,\ j = 1,2)$. Parameter estimates are obtained through the first step of the estimation procedure set forth in Section 3.1, relying on Indirect Inference and Simulated Method of Moments. The sample covers monthly data for the period from January 1950 to December 2006.

Parameter	Estimate	Std. error
κ_1	0.0231	0.0095
μ_1	1.0375	0.3784
α_1	2.4408×10^{-4}	1.0918×10^{-4}
β_1	-1.0005×10^{-6}	0.3374×10^{-6}
κ_2	0.9025	0.4037
μ_2	1.0386	0.3962
α_2	0.0253	0.0126
β_2	−0.0198	0.0084
$\overline{\kappa}_1$	−0.2995	0.1293
$\overline{\kappa}_2$	1.3723	0.6423

Table 2
Parameter estimates and block-bootstrap standard errors for stochastic secular growth G in Eq. (4), the real stock price in Eq. (8), and the unobservable factor $y_3(t)$ in Eq. (1). The parameter vector to be estimated is $\boldsymbol{\theta}^\top = (\boldsymbol{\theta}_u, (s_j)_{j=0}^3)$, where $\boldsymbol{\theta}_u = (\kappa_3, \mu_3, \alpha_3, \beta_3)$, and the long run mean of the unobservable factor, μ_3, is set equal to one for the purpose of model's identification. Parameter estimates are obtained through the second step of the estimation procedure set forth in Section 3.2, relying on Indirect Inference and Simulated Method of Moments, with parameters $\boldsymbol{\theta}_G^\top = (g, \sigma_G^2)$ estimated on the low frequency component of secular growth of the real stock price, extracted through the Hodrick–Prescott filter with smoothing parameter 1600. The sample covers monthly data for the period from January 1950 to December 2006.

Parameter	Estimate	Std. error
g	0.0264	0.0171
σ_G^2	0.0012	5.2089×10^{-4}
s_0	0.2272	0.1154
s_1	−0.8956	0.4426
s_2	1.8925	0.8873
s_3	−0.0560	0.0292
κ_3	0.0101	0.0043
μ_3	1	Restricted
α_3	1.2009	0.4699
β_3	0.0201	0.0106

factor and secular growth. Section 4.2.4 contains a descriptions of the unobserved factor dynamics implied by the model's estimates. Section 4.2.5 contains estimates of the model-implied volatility risk-premiums. Out-of-sample experiments relating to the subprime crisis are in Section 4.2.6.

4.2.1. Macroeconomic drivers

Table 1 reports parameter estimates and block-bootstrap standard errors for the joint process of the two macroeconomic variables, $y_{1,t}$ and $y_{2,t}$, as set forth in Section 3.1. The estimates are all largely significant, and confirm our discussion of Fig. 1: inflation is more persistent than IP growth, as both its speed of adjustment in the absence of feedbacks, κ_1, and its feedback parameter, $\overline{\kappa}_1$, are much lower than the counterparts for IP growth, κ_2 and $\overline{\kappa}_2$. Finally, the estimates of β_1 and β_2 are both negative, implying that the volatility of these two macroeconomic variables are countercyclical, an interesting property, from an asset pricing perspective. However, we note that the estimate of β_1, albeit statistically significant, is also economically very small.[5]

4.2.2. Aggregate stock returns and volatility

Table 2 reports estimates and block-bootstrap standard errors for (i) the parameters affecting secular growth, (ii) the parameters linking the two macroeconomic factors and the unobservable factor to the asset price, and (iii) the parameters for

[5] It is known since at least Friedman (1977) that high variability of inflation might link to high inflation. For example, Engle (1982) finds that inflation volatility increases during the middle 1970s. We have constructed measures of inflation volatility similar to that in Eq. (21), relating to the first difference of inflation, which confirm these findings. We also find that after the 1970s, inflation slowdowns tend to occur more rapidly than inflation increases although overall, a clear relation between inflation and inflation volatility is hard to establish. The estimate of β_1 for our continuous time model is likely to reflect these facts.

the unobservable factor process, as explained in Section 3.1. The estimates are all largely significant and point to two conclusions. First, the stock price is positively related to IP growth and negatively related to inflation. Second, the unobservable factor is quite persistent, displaying high volatility, as the estimate of the speed of mean reversion, κ_3, is low. Note, the literature on long run risks started by Bansal and Yaron (2004) emphasizes the asset pricing implications of long-run risks affecting the expected consumption growth rate. Interestingly, the presence of a persistent factor affecting stock returns and volatility emerges quite neatly from our estimation. However, in long-run risk models, expected consumption growth is unlikely to affect the dynamics of stock volatility which, instead, are inherited by those of the volatility of consumption growth. In our model, our unobservable factor does, instead, affect stock volatility, and substantially, as explained below.

Fig. 2 shows the dynamics of stock returns and volatility predicted by the model, along with their sample counterparts, calculated as described in Section 3.2. These predictions are obtained by feeding the model with sample data for the two macroeconomic factors, $y_{1,t}$ and $y_{2,t}$, in conjunction with simulations of the third unobservable factor, using all the estimated parameters. For each point in time, we average over the cross-section of 1000 simulations, and report returns (in the top panel of Fig. 2) and volatility (in the bottom panel). Returns are computed as we do with the data, and volatility is obtained through Eq. (12).

The model appears to capture the procyclical nature of stock returns and the countercyclical behavior of stock volatility. It generates *all* the market drops as well as *all* the volatility upward swings occurred during the NBER recessions, including the dramatic spike of the 1975 recession. In the data, average stock volatility is about 11.50%, with a standard deviation of about 4.0%. The model predicts an average volatility of about 13%, with a standard deviation of about 3.1%.

4.2.3. Factor decompositions

How much of the variation in volatility can be attributable to macroeconomic factors? It is a natural question, as the key innovation of our model is the introduction of these factors for the purpose of explaining volatility, on top of a standard unobservable factor. We address this issue and calculate (i) the ratio of the instantaneous return variance due to factor y_j, $s_j^2(\alpha_j+\beta_j y_j(t))$, to the total instantaneous variance, $\sigma^2(t)$ in Eq. (12) and (ii) the ratio of the instantaneous variance of secular growth, σ_G^2, to $\sigma^2(t)$, as follows:

$$C_j(t)s^2(\boldsymbol{y}(t)) \equiv \frac{s_j^2(\alpha_j+\beta_j y_j(t))}{\sigma^2(t)}, \quad j=1,2,3 \quad \text{and} \quad C_G(t) \equiv \frac{\sigma_G^2}{\sigma^2(t)}, \tag{27}$$

where $s(\boldsymbol{y}) \equiv s_0 + \sum_{j=1}^{3} s_j y_j$. Fig. 3 depicts the time series of $C_j(t)$ and $C_G(t)$ implied by our estimated model, obtained, as usual, by feeding the model with the observed samples of $y_1(t)$ and $y_2(t)$, and averaging across 1000 simulations of $y_3(t)$. The clear finding is that industrial production growth makes the most important contribution to stock volatility: the time series average of $C_2(t)$ is above 73%, more than four times higher than that of $C_3(t)$, the contribution made by the unobserved factor. Panel A of Table 3 reports averages and standard deviations of the contributions made by all factors, and secular growth. Variations in industrial production growth and the unobserved factor are responsible, alone, for more than 90% of the variation in stock volatility. It is a striking result, as one challenge we face is to explain why we have observed a sustained stock market volatility, in spite of the Great Moderation. Our estimated model entails two clear conclusions.

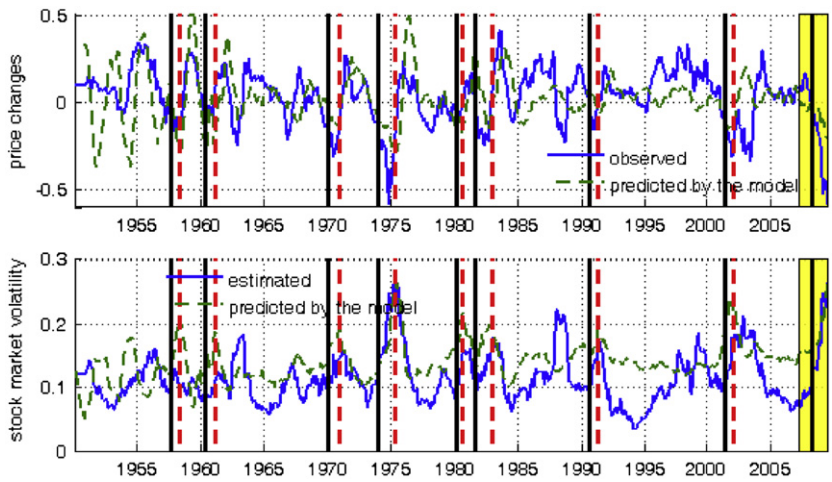

Fig. 2. This figure plots one-year ex post price changes, defined as $R_t \equiv \ln s_t/s_{t-12}$ and one-year return volatility, defined as $\mathrm{Vol}_t \equiv \sqrt{r\pi}(1/12)\sum_{i=t}^{1} 2|\ln s_{t+1-i}/s_{t-i}|$, along with their counterparts predicted by the model. The instantaneous standard deviation predicted by the model is obtained through Eq. (12). The sample covers monthly data for the period from January 1950 to December 2006. Vertical solid lines track the beginning of NBER-dated recessions and vertical dashed lines indicate the end of NBER-dated recessions. The shaded area covers the out-of-sample period, from January 2007 to March 2009.

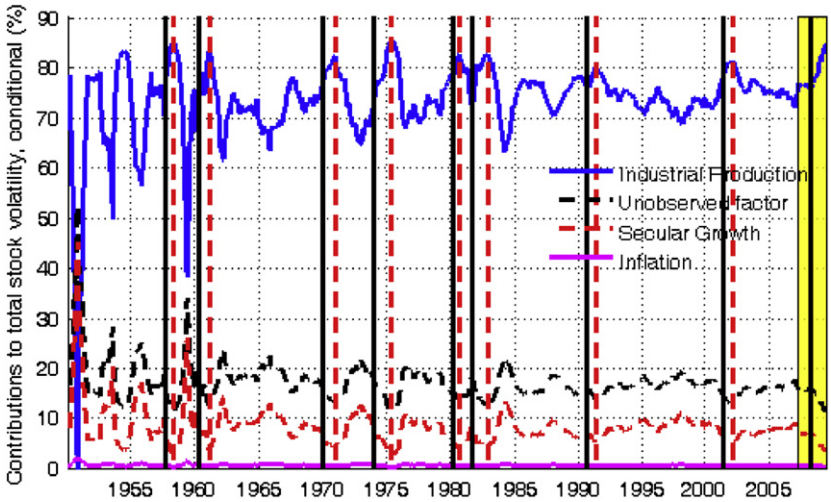

Fig. 3. This figure plots the contributions to stock volatility made by macroeconomic and unobservable factors, and that made by secular growth, as defined by C_j, $j=1$, 2, 3 and C_G in Eq. (27). From top to bottom, "Industrial Production" is C_2, "Unobservable factor" is C_3, "Secular Growth" is C_G, and "Inflation" is C_1. The sample covers monthly data for the period from January 1950 to December 2006. Vertical solid lines track the beginning of NBER-dated recessions and vertical dashed lines indicate the end of NBER-dated recessions. The shaded area covers the out-of-sample period, from January 2007 to March 2009.

Table 3
Variance decomposition statistics for (i) realized volatility, as calculated in Section 4.2.2 (Panel A) and expected volatility under the risk-neutral probability, as calculated in Section 4.2.5 (Panel B). Panel A reports averages and standard deviations of the contributions $C_j(t)$ and $C_G(t)$ to the total variance, $\sigma^2(t)$ in Eq. (12), made by: (i) the two macroeconomic factors, gross inflation, $y_1(t)$, and gross industrial production growth, $y_2(t)$, as defined in Table 1, (ii) the unobserved factor, $y_3(t)$, and (iii) secular growth, defined as in Eq. (27). The sample covers monthly data for the period from January 1950 to December 2006. Panel B reports statistics for the risk-neutral counterparts to the average paths of $C_j(t)$ and $C_G(t)$. The sample covers monthly data for the period from January 1990 to December 2006.

Panel A: Contributions of factors to stock volatility

	1950–2006 (%)	1950–1981 (%)	1960–1981 (%)	1982–2006 (%)
Averages				
Gross inflation	0.87	0.92	0.88	0.83
Gross growth	73.47	71.97	73.57	75.10
Unobserved factor	17.23	18.10	17.41	16.27
Secular growth	8.43	9.01	8.13	7.78
Standard deviations				
Gross inflation	0.18	0.23	0.12	0.08
Gross growth	7.53	9.65	5.03	3.47
Unobserved factor	3.69	4.67	2.46	1.71
Secular growth	3.67	4.75	2.40	1.68

Panel B: Contributions of factors to the VIX index

	Averages (%)	Standard deviations (%)
Gross inflation	1.55	0.03
Gross growth	86.91	1.28
Unobserved factor	8.79	0.56
Secular growth	2.75	0.17

First, as Fig. 3 makes clear, the 73% average contribution of industrial production growth to stock volatility seems to be rather stable over time, at least once we exclude the 1950s—a period of sustained volatility for growth (see Fig. 1). Accordingly, the Great Moderation does merely appear to have affected the variability of the linkages between industrial production growth and aggregate stock volatility, not the very same linkages. To illustrate, Panel A of Table 3 shows averages and standard deviations of the factors' contributions across different sampling periods. We take 1982 to be the year that marks the beginning of the Great Moderation, characterized by the inauguration of the Federal Reserve monetary policy turning point and a lower volatility of real macroeconomic variables (e.g., Blanchard and Simon, 2001). As is clear, whilst the average contributions are stable, the variability of these contributions has decreased over the Great Moderation.

For example, the average of $C_2(t)$ is between 73% and 75%, across all sampling periods, whereas its standard deviation decreases to 3.47% during the 1982–2006 sample, from 9.65% (1950–1981) and 5.03% (1960–1981).

Second, the contribution of industrial production growth to volatility, albeit crucial, is not exhaustive. Our model predicts that stock volatility cannot be explained by macroeconomic variables only, as the unobserved factor accounts for about 17% of the fluctuations in $\sigma^2(t)$. Equally important is the observation that the contribution of industrial production to stock volatility is strongly countercyclical, exhibiting large upward swings starting at, and sometimes, anticipating, turning points, as in the case of the 1970 recessions and the most recent, 2001 recession. Instead, the contribution of the unobserved factor to stock volatility, $C_3(t)$, is procyclical, for the simple reason that the instantaneous volatility of $y_3(t)$ does not obviously link to the business cycle, thereby making the ratio $C_3(t)$ in Eq. (27) procyclical, due to the countercyclical nature of its denominator, $\sigma^2(t)$. Altogether, our empirical results suggest that while unobserved factors are needed to explain the level of stock volatility, industrial production is needed to explain the countercyclical swings of stock volatility that we have in the data—the volatility of volatility.

Finally, the contribution of secular growth to stock volatility is limited, being approximately 8%, and that of gross inflation plays an even more marginal role, being less than 1%. Note, however, that our model predicts that inflation links to asset returns and volatility in a manner comparable to that in the data. For example, it is well-known since at least Fama (1981) that real stock returns are negatively correlated with inflation, a property that hinders the ability of stocks to hedge against inflation. In our sample, this correlation is -35%, while the correlation our model generates is -24%. Finally, the correlation between stock volatility and inflation is about 20% in the data, while that implied by the model is about 25%.

4.2.4. Model-implied unobserved factor dynamics

The predictions of the model discussed so far rely on cross-sectional averages of simulations of the unobserved factor, y_3. Yet what is the interpretation of this unobserved factor? Let us invert the price function in Eq. (8), for y_3, and for each month, as follows:

$$-\hat{y}_{3,t} \equiv -\frac{1}{\hat{s}_3}\left(\frac{S_t}{\hat{G}_t} - \hat{s}_0 - \hat{s}_1 y_{1,t} - \hat{s}_2 y_{2,t}\right), \tag{28}$$

where S_t is the real stock price at time t, $(\hat{s}_j)_{j=0}^{3}$ are estimates of the pricing function coefficients, as reported in Table 2, and $\hat{G}_t$ is the cross-sectional average of 1000 simulations of secular growth.

Fig. 4 (top panel) depicts $-\hat{y}_{3,t}$ (in bold), along with 100 simulated trajectories of the unobserved factor performed with the parameter estimates in Table 2. Reinsuringly, the range of variation of the model-implied factor roughly falls within that of the simulated trajectories of this factor. Note that the estimate of s_3 is negative, such that $-\hat{y}_{3,t}$ positively affects the real stock price—higher realizations of $-\hat{y}_{3,t}$ amount to good pieces of news to the stock market. There are episodes where $-\hat{y}_{3,t}$ comes close to the edges of the realized range of variation experienced by the unobserved factors during the simulations. These episodes are interesting, as they correspond to: (i) the lows of the late 1970s and the early 1980s, and (ii) the highs of the dotcom bubble of the late 1990s. The extracted factor oscillates between (about) its minimum and its maximum over those approximate twenty years. The rise and fall over this period have a clear economic interpretation,

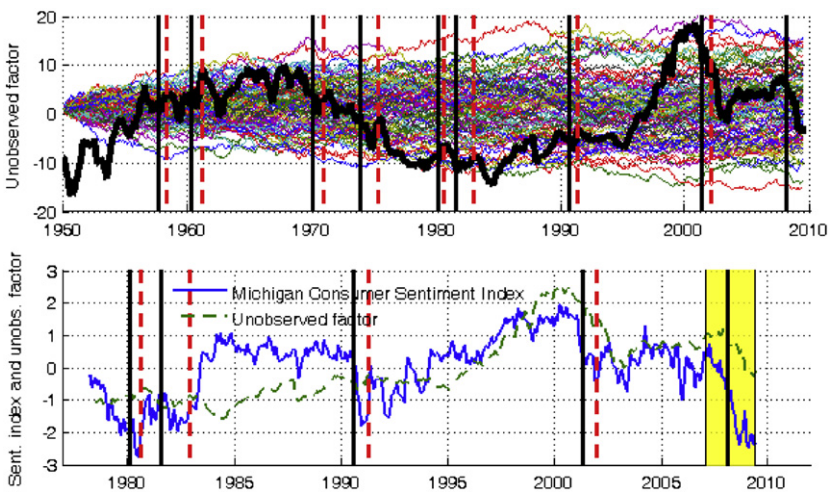

Fig. 4. This figure plots estimates of the model-implied unobserved factor, and a comparison of it with the University of Michigan Consumer Sentiment (UMCSENT) index. The top panel plots: (i) 100 simulated trajectories of the unobserved factor and (ii) (minus) the model-implied unobserved factor (the bold line), as defined in Eq. (28). The bottom panel depicts the UMCSENT index, de-meaned and standardized by its own standard deviation, together with the negative of the model-implied unobserved factor, de-meaned and standardized by its own standard deviation. The sample covers monthly data for the period from January 1950 to December 2006. The UMCSENT index is available for the sample period starting from January 1978. Vertical solid lines track the beginning of NBER-dated recessions and vertical dashed lines indicate the end of NBER-dated recessions. The shaded area covers the out-of-sample period, from January 2007 to March 2009.

with the late 1970s and early 1980s being particularly bad times, marked by the occurrence of a double dip recession, and the extraordinary market boom over the dotcom bubble being notoriously suspected to be one of exuberance (e.g., Shiller, 2005). These observations motivate us to explore the extent to which our extracted factor links to indexes of "sentiment," following a recent strand of the literature that attempts to link asset price movements to factors such as investors uncertainty (as in David and Veronesi, 2006), confidence risk (as in Bansal and Shaliastovich, 2010), or Knightian uncertainty (as in Drechsler, 2010, or Mele and Sangiorgi, 2011). The bottom panel of Fig. 4 compares the time series behavior of the model-implied unobserved factor, $-\hat{y}_{3,t}$, with that of an index of consumer confidence—the University of Michigan Consumer Sentiment (UMCSENT) index, available from January 1978.

Note how the UMCSENT index tracks the lows and the highs of the market that have so slowly occurred over the last 30 years: the bad times of the late 1970s and early 1980s, the rise occurring over the late 1980s and culminating with the dotcom bubble of the late 1990s and, finally, the drop of the late 2000s, corresponding to the subprime crisis—a period analyzed in more detail in the next section. Interestingly, our extracted factor, $-\hat{y}_{3,t}$, co-moves positively with the UMCSENT index, correlating with it at about 50%. In contrast, its correlation with the macroeconomic factors is modest (10% with inflation and 30% with industrial production growth). Interestingly, then, the pattern our extracted factor exhibits is one that mostly tracks long-run movements of the market, even more so than the short-term movements relating to business cycles. In Appendix C.3 of the Supplementary material, we produce variance decompositions statistics, obtained by feeding the model with both $\hat{y}_{3,t}$ in Eq. (28) and the UMCSENT index, instead of relying on simulations of y_3, and document results similar to those summarized by Fig. 3 and Table 3.

4.2.5. Volatility risk-premiums and the dynamics of the VIX index

Table 4 reports parameter estimates and block-bootstrap standard errors for the vector of the risk-premiums coefficients $\boldsymbol{\lambda}$ in Eq. (7), as set forth in Section 3.3. The estimates, all significant, imply that the risk-premiums processes are all positive, and quite large, especially those relating to the two macroeconomic factors. Moreover, the risk compensation for inflation increases with inflation and that for industrial production is countercyclical, given the sign of the estimated values for the loadings of inflation, $(\beta_1\lambda_{1(1)}+\lambda_{2(1)})$ (positive), and industrial production, $(\beta_2\lambda_{1(2)}+\lambda_{2(2)})$ (negative), in the risk-premium process of Eq. (7). While gross inflation does receive compensation, the countercyclical behavior of the risk-premium for industrial production growth is even more critical, as we explain below. Our estimated model predicts that in bad times, the risk-premium for industrial production growth goes up, and future expected economic conditions even worsen, under the risk-neutral probability, which boosts future expected volatility, under the same risk-neutral probability. In part because of these effects, the VIX index predicted by the model is countercyclical. This reasoning is quantitatively sound. Fig. 5 (top panel) depicts the VIX index, along with the VIX index predicted by the model and the (square root of the) model-implied expected integrated variance, obtained through simulations performed as explained in previous sections. The model appears to track the swings the VIX index has undergone over the 1991 and the 2001 recession episodes.

The top panel of Fig. 5 also shows the dynamics of volatility expected under the physical probability. This expected volatility is certainly countercyclical, although it does not display the large variations the model predicts for its risk-neutral counterpart, the VIX index. The VIX index predicted by the model is countercyclical because, as explained, the risk-premiums required to bear the fluctuations of the macroeconomic factors are (i) positive and (ii) countercyclical, and, also, because (iii) current volatility is countercyclical. Under the physical probability, expected volatility is countercyclical only because of the third effect. However, quantitatively, movements of volatility risk-premiums account for variations in the VIX index sensibly more than those of the volatility expected under the physical probability, as clearly summarized by Fig. 5.

Table 4
Parameter estimates and block-bootstrap standard errors for the risk-premium parameters of the total risk-premium process in Eq. (7). The parameter vector to be estimated is $\boldsymbol{\lambda}^\top=(\lambda_{1(1)},\lambda_{2(1)},\lambda_{1(2)},\lambda_{2(2)},\lambda_{1(3)},\lambda_{2(3)})$. Parameter estimates are obtained through the third step of the estimation procedure set forth in Section 3.3, relying on Indirect Inference and Simulated Method of Moments. The sample covers monthly data for the period from January 1990 to December 2006.

	Estimate	Std. error
Inflation		
$\lambda_{1(1)}$	-2.1533×10^3	0.9683×10^3
$\lambda_{2(1)}$	32.0141	15.5655
Ind. prod.		
$\lambda_{1(2)}$	5.6760×10^2	2.7643×10^2
$\lambda_{2(2)}$	5.5717	2.7952
Unobs.		
$\lambda_{1(3)}$	0.0019	0.0008
$\lambda 2(3)$	5.9837×10^{-4}	2.9526×10^{-4}

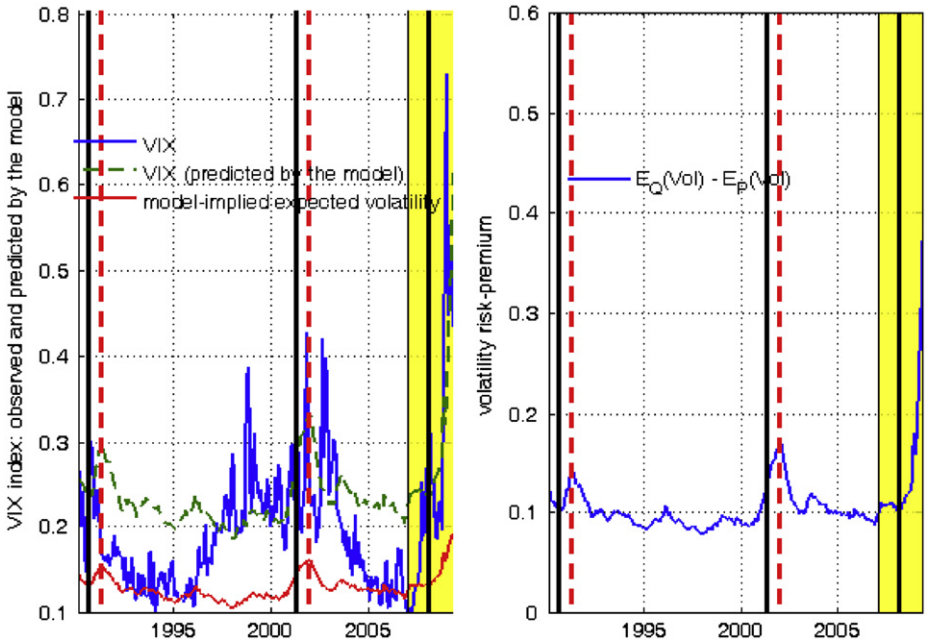

Fig. 5. This figure plots the VIX index along with model's predictions. The top panel depicts (i) the VIX index, (ii) the VIX index predicted by the model, and (iii) the VIX index predicted by the model in an economy without risk-aversion, i.e. the expected integrated volatility under the physical probability. The bottom panel depicts the volatility risk-premium predicted by the model, defined as the difference between the model-generated expected integrated volatility under the risk-neutral and the physical probability, as in Eq. (16). The sample covers monthly data for the period from January 1990 to December 2006. Vertical solid lines track the beginning of NBER-dated recessions and vertical dashed lines indicate the end of NBER-dated recessions. The shaded area covers the out-of-sample period, from January 2007 to March 2009.

Which factors mostly contribute to the dynamics of the VIX? Panel B of Table 3 reports averages and standard deviations of the contributions of each factor, as predicted by our estimated model. We calculate each of these contributions by evaluating C_j and C_G in Eq. (27) under the risk-neutral probability and, then, aggregating the average paths of C_j and C_G for every month, and, finally, taking cross-sectional averages over 1000 simulations of the unobserved factor. Similarly as for the results in Section 4.2.2 on realized volatility, we find, again, that our model predicts industrial production growth to account for the bulk of variation of the VIX index. The unobserved factor accounts for less than 10%, and inflation and secular growth play a quite marginal role, explaining no more than 5%, of the model-implied VIX. Interestingly, Stock and Watson (2003) find that the linkages of asset prices to growth are stronger than for inflation. Our results further qualify this finding: inflation does not seem to affect too much the dynamics of neither realized volatility nor future expected volatility under the risk-neutral probability.

Finally, the bottom panel in Fig. 5 plots the volatility risk-premium, defined as in Eq. (16). This risk-premium is countercyclical, and this property arises for exactly the same reasons we put forward to explain the swings the model predicts for the VIX index: positive compensation for risk, combined with countercyclical variation of the premiums required to compensate for the risk of fluctuations of the macroeconomic factors.

4.2.6. Out-of-sample predictions, and the subprime crisis

We undertake out-of-sample experiments to investigate the model's predictions over a quite exceptional period, that from January 2007 to March 2009. This sample covers the subprime turmoil, and features unprecedented events, both for the severity of capital markets uncertainty and the performance of the US economy. The market witnessed to a spectacular drop accompanied by an extraordinary surge in volatility. In March 2009, yearly returns plummeted to -58.30%, a performance even worse than that experienced in October 1974 (-58.10%). Furthermore, according to our estimates, obtained through Eq. (21), aggregate stock volatility reached 28.20% in March 2009, the highest level ever experienced in our sample. Finally, the VIX index hit its highest value in our sample in November 2008 (72.67%), and remained stubbornly high for several months. The time series behavior of stock returns, stock volatility and the VIX index during our out-of-sample period are depicted over the shaded areas in Figs. 2 and 5.

Macroeconomic developments over our out-of-sample period (the shaded area in Fig. 1) were equally extreme, with yearly inflation rates achieving negative values in 2009, and yearly industrial production growth being as low as -13%, in March 2009. Under such macroeconomic conditions, we expect our model to produce the following predictions: (i) stock returns drop, (ii) stock volatility rises, (iii) the VIX index rises, and more than stock volatility. The mechanism is, by now, clear. Asset prices and, hence, returns, plummet, as they are positively related to growth, which crashed. (Note that inflation also decreased but the (negative) price sensitivity to it is much smaller than that of growth.) Moreover, volatility

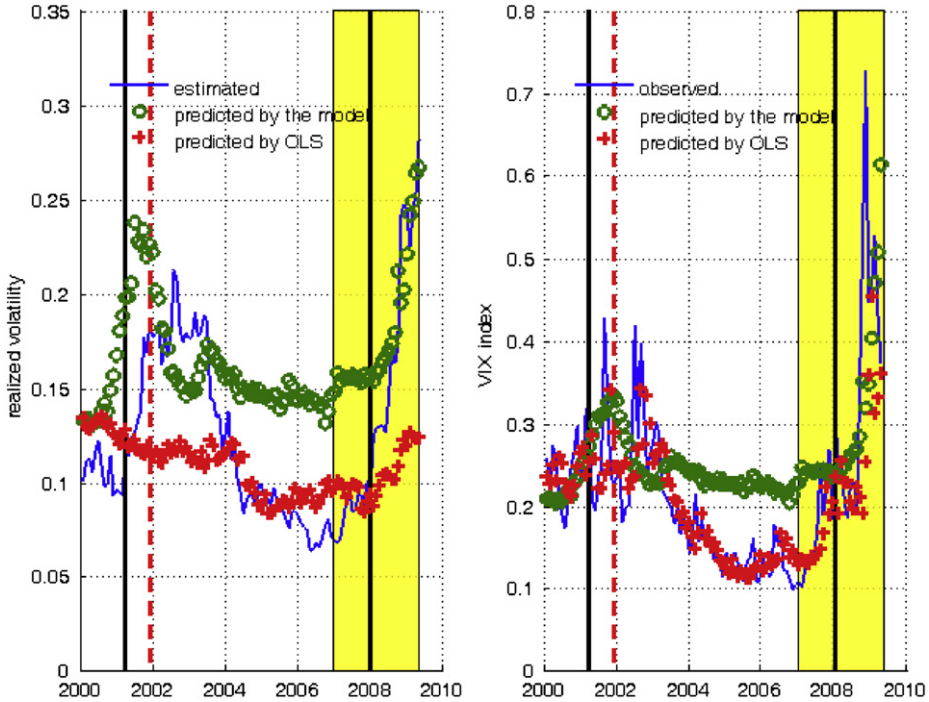

Fig. 6. The left panel of this figure plots the return volatility Vol_t, as defined in Fig. 2, along with the instantaneous standard deviation predicted by (i) the model, through Eq. (12) and (ii) the predictive part of an OLS regression of Vol_t on to past values of Vol_t, inflation and industrial production growth. The right panel depicts the VIX index, along with the VIX index predicted by (i) the model and (ii) the predictive part of an OLS regression of the VIX index on to past values of the VIX index, inflation and industrial production growth. The sample period is from January 2000 to March 2009, and the estimation of the model and OLS regressions relates to the period from January 1950 to December 2006. Vertical solid lines track the beginning of NBER-dated recessions and the vertical dashed line indicates the end of the recession occurred in November 2001. The shaded area covers the out-of-sample period from January 2007 to March 2009, which includes the NBER recession announced to have occurred in December 2007, and the subprime crisis, which started in June 2007.

increases, with the VIX index increasing even more, due to our previous finding of (i) sizeable macroeconomic risk-premiums and (ii) strong countercyclical variation in these premiums.

Figs. 2 and 5 confirm our reasoning, and reveal that the model is able to trace out the dynamics of stock returns and volatility (Fig. 2), and the VIX index (Fig. 5), over the out-of-sample period. The market literally crashes, as in the data, although only less than a half as much as in the data: the lowest value for yearly stock returns the model predicts, out-of-sample, is -21.77%, which is the second lowest figure our model produces, since after the quite volatile periods occurring over the 1950s and the early 1960s. (The lowest level the model predicts after those periods is -29.91%, for March 1975, the last month of the second severe recession of the 1970s.) Instead, the model predicts that stock volatility and the VIX index surge even more than in the data, reaching record highs of 26.68% (volatility) and 61.27% (VIX).

Fig. 6 provide additional details regarding the period from January 2000 to March 2009. It compares stock volatility and the VIX index with the predictions of the model and those of an OLS regression. The OLS for volatility is that in Eq. (23), excluding the lag for six months, related to the autoregressive term. The OLS for the VIX index is that in Eq. (25). OLS predictions are obtained by feeding the OLS predictive part with its regressors, using parameter estimates obtained with data up to December 2006. The following table reports Root Mean Squared Errors (RMSE) for both our model and OLS, calculated over the out-of-sample period.

RMSE for the model and OLS

	Model	OLS
Volatility	0.0478	0.0700
VIX index	0.1119	0.1319

Overall, OLS predictions do not seem to capture the countercyclical behavior of stock volatility. As for the VIX index, the OLS model (in fact, by Eq. (25), an autoregressive, distributed lag model) produces predictions that are not as accurate as the model, and generate overfit. The model, instead, predicts the swings we see in the data, in both the last two recession episodes. The RMSEs clearly favour the model against OLS, although it appears to do so more with realized volatility than with the VIX index, as Fig. 6 informally reveals. Appendix C.3 of the Supplementary material confirms these findings in

additional experiments, performed by feeding the model with both $\hat{y}_{3,t}$ in Eq. (28) and the UMCSENT index, rather than by simulations of y_3.

5. Conclusion

How does aggregate stock market volatility relate to the business cycle? This old question has been formulated at least since Officer (1973) and Schwert (1989a,b). We learnt from recent theoretical explanations that the countercyclical behavior of stock volatility can be understood as the result of a rational valuation process. However, how much of this countercyclical behavior is responsible for the sustained level aggregate volatility has experienced for centuries? This paper develops a model where approximately 75% of the variations in stock volatility can be explained by macroeconomic factors, and where some unobserved component is also needed to make stock volatility consistent with rational asset valuation.

Our model predicts that risk-premiums arising from fluctuations in this volatility are strongly countercyclical, certainly more so than stock volatility alone. In fact, the risk-compensation for the fluctuation of the macroeconomic factors is large and countercyclical, and helps explain the swings in the VIX index that we observe during recessions. In out-of-sample experiments that cover the 2007–2009 subprime crisis, the VIX reached a record high of more than 70%, which our model can at least partially track, through a countercyclical variation in the volatility risk-premiums. Again, our model predicts that a business cycle factor such as industrial production growth can explain more than 85% of the variations of the VIX index.

The key aspect of our model is that the relations among the market, stock volatility, volatility risk-premiums and the macroeconomic factors, are consistent with no-arbitrage. In particular, volatility is endogenous in our framework: the same variables driving the payoff process and the volatility of the pricing kernel, and hence, the asset price, are those that drive stock volatility and volatility-related risk-premiums. A question for future research is to explore whether the no-arbitrage framework in this paper can be used to improve forecasts of real economic activity. In fact, stock volatility and volatility risk-premiums are driven by business cycle factors, as this paper clearly demonstrates. A challenging and fundamental question is to explore the extent to which business cycle, stock volatility and volatility risk-premiums do endogenously develop.

Acknowledgments

We wish to thank Urban Jermann (the co-editor), two anonymous referees, and Yacine Aït-Sahalia, Alessandro Beber, Bernard Dumas, Marcelo Fernandes, Daniel Gutknecht, Christian Julliard, Nour Meddahi, Mark Salmon, George Tauchen, Viktor Todorov, and seminar participants at EDHEC-Nice, European Central Bank, HEC Lausanne, Hong Kong Monetary Authority, IE Business School (Madrid), the Universities of Aarhus (CREATES), Bicocca (Milan), Goethe (Frankfurt), Louvain, Lugano (Finance) and Warwick, the 2006 London-Oxford Financial Econometrics Study Group, the 2006 conference on Realized Volatility (Montréal), the 2008 UCL-CEMMAP conference, the 2008 Imperial College Financial Econometrics Conference, the 2008 Adam Smith Asset Pricing Conference at London Business School, the 2008 LSE-FMG conference on "Integrating historical data and market expectations in financial econometrics," the 2008 North American Summer Meeting of the Econometric Society (Carnegie), and the 2008 SITE Summer Workshop at Stanford University for valuable comments and suggestions. Valentina Corradi and Walter Distaso acknowledge support from the British ESRC under the grant RES-062-23-0311. Antonio Mele thanks the British EPSRC for support via grant EP/C522958/1. Part of this research has been carried out within the project on "New Methods in Theoretical and Empirical Asset Pricing" of the National Centre of Competence in Research "Financial Valuation and Risk Management" (NCCR FINRISK). The NCCR FINRISK is a research instrument of the Swiss National Science Foundation. The usual disclaimer applies.

Appendix A. Supplementary material

Supplementary data associated with this article can be found in the online version at http://dx.doi.org.10.1016/j.jmoneco.2012.10.019.

References

Adrian, T., Rosenberg, J., 2008. Stock returns and volatility: pricing the short-run and long-run components of market risk. Journal of Finance 63, 2997–3030.

Aït-Sahalia, Y., 2008. Closed-form likelihood expansions for multivariate diffusions. Annals of Statistics 36, 906–937.

Aït-Sahalia, Y., Mykland, P.A., Zhang, L., 2011. Ultra high frequency volatility estimation with dependent microstructure noise. Journal of Econometrics 160, 160–175.

Altissimo, F., Mele, A., 2009. Simulated nonparametric estimation of dynamic models. Review of Economic Studies 76, 413–450.

Ang, A., Piazzesi, M., 2003. A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables. Journal of Monetary Economics 50, 745–787.

Ang, A., Piazzesi, M., Wei, M., 2006. What does the yield curve tell us about GDP growth? Journal of Econometrics 131, 359–403.

Bakshi, G.S., Chen, Z., 1996. Inflation, asset prices, and the term structure of interest rates in monetary economies. Review of Financial Studies 9, 241–275.
Bakshi, G., Madan, D., 2000. Spanning and derivative security evaluation. Journal of Financial Economics 55, 205–238.
Bakshi, G., Madan, D., 2006. A theory of volatility spreads. Management Science 52, 1945–1956.
Bansal, R., Yaron, A., 2004. Risks for the long run: a potential resolution of asset pricing puzzles. Journal of Finance 59, 1481–1509.
Bansal, R., Shaliastovich, I., 2010. Confidence risk and asset prices. American Economic Review P&P 100, 537–541.
Barndorff-Nielsen, O.E., Hansen, P.R., Lunde, A., Shephard, N., 2008. Designing realized kernels to measure the ex-post variation of equity prices in the presence of noise. Econometrica 76, 1481–1536.
Bernanke, B.S., 2004. The great moderation: remarks by the Federal Reserve Board Governor, The Meetings of the Eastern Economic Association, Washington, DC, February 20.
Blanchard, O., Simon, J., 2001. The long and large decline in U.S. output volatility. Brookings Papers on Economic Activity 1, 135–164.
Bloom, N., 2009. The impact of uncertainty shocks. Econometrica 77, 623–685.
Bloom, N., Floetotto, M., Jaimovich, N., 2009. Really Uncertain Business Cycles. Working Paper, Stanford University.
Bollerslev, T., Zhou, H., 2006. Volatility puzzles: a simple framework for gauging return-volatility regressions. Journal of Econometrics 131, 123–150.
Bollerslev, T., Gibson, M., Zhou, H., 2011. Dynamic estimation of volatility risk premia and investor risk aversion from option implied and realized volatilities. Journal of Econometrics 160, 102–118.
Bollerslev, T., Tauchen, G., Zhou, H., 2009. Expected stock returns and variance risk premia. Review of Financial Studies 22, 4463–4492.
Brandt, M.W., Kang, Q., 2004. On the relationship between the conditional mean and volatility of stock returns: a latent VAR approach. Journal of Financial Economics 72, 217–257.
Britten-Jones, M., Neuberger, A., 2000. Option prices, implied price processes and stochastic volatility. Journal of Finance 55, 839–866.
Buraschi, A., Jiltsov, J., 2005. Inflation risk premia and the expectations hypothesis. Journal of Financial Economics 75, 429–490.
Campbell, J.Y., 2003. Consumption-based asset pricing. In: Constantinides, G.M., Harris, M., Stulz, R.M. (Eds.), Handbook of the Economics of Finance, vol. 1B. North-Holland Elsevier, Amsterdam, pp. 803–887. (chapter 13).
Campbell, J.Y., Cochrane, J.H., 1999. By force of habit: a consumption-based explanation of aggregate stock market behavior. Journal of Political Economy 107, 205–251.
Campbell, J.Y., Lettau, M., Malkiel, B.G., Xu, Y., 2001. Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk. Journal of Finance 56, 1–43.
Carr, P., Madan, D., 2001. Optimal positioning in derivative decurities. Quantitative Finance 1, 19–37.
Carr, P., Wu, L., 2009. Variance risk premiums. Review of Financial Studies 22, 1311–1341.
Carrasco, M., Chernov, M., Florens, J.-P., Ghysels, E., 2007. Efficient estimation of general dynamic models with a continuum of moment conditions. Journal of Econometrics 140, 529–573.
Chen, N.-F., Roll, R., Ross, S.A., 1986. Economic forces and the stock market. Journal of Business 59, 383–403.
Chernov, M., Ghysels, E., 2000. A study towards a unified approach to the joint estimation of objective and risk-neutral measures for the purpose of options valuation. Journal of Financial Economics 56, 407–458.
Cochrane, J.H., 2005. Financial Markets and the Real Economy: Foundations and Trends® in Finance. Now Publishers Inc., Hanover, MA.
Corradi, V., Distaso, W., 2006. Semiparametric comparison of stochastic volatility models using realized measures. Review of Economic Studies 73, 635–667.
Danthine, J.-P., Donaldson, J.B., 1986. Inflation and asset prices in an exchange economy. Econometrica 54, 585–605.
David, A., Veronesi, P., 2006. Inflation and Earnings Uncertainty and Volatility Forecasts. Working Paper, University of Calgary and University of Chicago.
Drechsler, I., 2010. Uncertainty, Time-Varying Fear, and Asset Prices. Working Paper, New York University.
Drechsler, I., Yaron, A., 2011. What's vol got to do with it. Review of Financial Studies 24, 1–45.
Duffee, G.R., 2002. Term premia and interest rate forecasts in affine models. Journal of Finance 57, 405–443.
Duffie, D., Singleton, K., 1993. Simulated moments estimation of Markov models of asset prices. Econometrica 61, 929–952.
Duffie, D., Kan, R., 1996. A yield-factor model of interest rates. Mathematical Finance 6, 379–406.
Dumas, B., 1995. The Meaning of the Implicit Volatility Function in case of Stochastic Volatility. Working Paper, HEC Lausanne. Available from: 〈http://www.insead.edu/facultyresearch/faculty/personal/bdumas/research/index.cfm〉.
Engle, R.F., 1982. Autoregressive conditional heteroskedasticity with estimates of the variance of the United Kingdom inflation. Econometrica 50, 987–1007.
Engle, R.F., Rangel, J.G., 2008. The spline GARCH model for low frequency volatility and its macroeconomic causes. Review of Financial Studies 21, 1187–1222.
Fama, E.F., 1981. Stock returns, real activity, inflation, and money. American Economic Review 71, 545–565.
Fermanian, J.-D., Salanié, B., 2004. A nonparametric simulated maximum likelihood estimation method. Econometric Theory 20, 701–734.
Fornari, F., Mele, A., 2010. Financial Volatility and Economic Activity. Working Paper, European Central Bank and London School of Economics.
Friedman, M., 1977. Nobel lecture: inflation and unemployment. Journal of Political Economy 85, 451–472.
Gallant, A.R., Tauchen, G., 1996. Which moments to match? Econometric Theory 12, 657–681.
Garcia, R., Lewis, M.-A., Pastorello, S., Renault, E., 2011. Estimation of objective and risk-neutral distributions based on moments of integrated volatility. Journal of Econometrics 160, 22–32.
Gouriéroux, C., Monfort, A., Renault, E., 1993. Indirect inference. Journal of Applied Econometrics 8, S85–S118.
Guvenen, F., 2009. A parsimonious macroeconomic model for asset pricing. Econometrica 77, 1711–1740.
Hamilton, J.D., Lin, G., 1996. Stock market volatility and the business cycle. Journal of Applied Econometrics 11, 573–593.
Heston, S.L., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. Review of Financial Studies 6, 327–343.
Hodrick, R.J., Prescott, E.C., 1997. Postwar U.S. business cycles: an empirical investigation. Journal of Money, Credit and Banking 29, 1–16.
King, M., Sentana, E., Wadhwani, S., 1994. Volatility and links between national stock markets. Econometrica 62, 901–933.
Mehra, R., Prescott, E.C., 2003. The equity premium in retrospect. In: Constantinides, G.M., Harris, M., Stulz, R.M. (Eds.), Handbook of the Economics of Finance, vol. 1B. North-Holland Elsevier, Amsterdam, pp. 889–938. (chapter 14).
Mele, A., 2007. Asymmetric stock market volatility and the cyclical behavior of expected returns. Journal of Financial Economics 86, 446–478.
Mele, A., Sangiorgi, F., 2011. Ambiguity, Information Acquisition and Price Swings in Asset Markets. Working Paper, London School of Economics and Stockholm School of Economics.
Officer, R.R., 1973. The variability of the market factor of the New York Stock Exchange. Journal of Business 46, 434–453.
Schwert, G.W., 1989a. Why does stock market volatility change over time? Journal of Finance 44, 1115–1153.
Schwert, G.W., 1989b. Business cycles, financial crises and stock volatility. In: Carnegie-Rochester Conference Series on Public Policy, vol. 31, 1989, pp. 83–125.
Shiller, R.J., 2005. Irrational Exuberance. Princeton University Press, Princeton.
Stock, J.H., Watson, M.W., 2003. Forecasting output and inflation: the role of asset prices. Journal of Economic Literature 41, 788–829.
Tauchen, G., 2005. Stochastic Volatility in General Equilibrium. Working Paper, Duke University.
Todorov, V., 2010. Variance risk premium dynamics: the role of jumps. Review of Financial Studies 23, 345–383.
Todorov, V., Tauchen, G., 2011. Volatility jumps. Journal of Business and Economic Statistics 29, 356–371.