Understanding Stock Market Volatility

A Business Cycle Perspective*

Antonio Mele
London School of Economics

April 21, 2008

Introduction

One of the most prominent features of the U.S. stock market is the close connection between aggregate stock market volatility and the development of the business cycle. Figure 1 depicts the statistical relation between stock market volatility and the industrial production growth rate over the last sixty years, which shows that stock volatility is largely countercyclical, being larger in bad times than in good times. Understanding the origins and implications of these facts is extremely relevant to policy makers. Indeed, if stock market volatility is countercyclical, it must necessarily be encoding information about the development of the business cycle. Policy makers could then attempt at extracting the signals stock volatility brings about the development of the business cycle.

In this short essay, I accomplish three tasks. First, I review and uncover a few stylized facts about stock market volatility, arising within a business cycle perspective (in Section 1). Second, I seek for theoretical explanations of these facts (in Section 2). Third, I investigate whether stock market volatility contains any useful information about the evolution of the business cycle (in Section 3). There are other exciting topics left over from this essay. For example, I do not tackle statistical issues related to volatility measurement (see, e.g., Andersen, Bollerslev and Diebold (2002) for a survey on

*This paper is a substantial revision of a previous article titled, "Understanding Stock Market Volatility," which appeared in the *Financial Markets Group Review, London School of Economics* 67, 10-15 (October 2005). In writing this essay, I have greatly benefited from the many discussions with my co-authors working on projects related to the topics dealt with in the essay: Valentina Corradi, Walter Distaso and Fabio Fornari. I am also indebted to Andrew Patton for his comments on the very first draft of this paper. I thank the British EPSRC for financial support via grant EP/C522958/1. The usual disclaimer applies.

the many available statistical techniques to estimate volatility). Nor do I consider the role of volatility in risk-management, portfolio selection, or derivative pricing (see, e.g., Lewis (2000) for a thorough analysis of these issues). At a more fundamental level, the focus of this essay is to explore the extent to which stock market volatility movements can be given a wider business cycle perspective, and to highlight some of the rational mechanisms underlying them.

1. Volatility cycles

Why is stock market volatility related to the business cycle? Financial economists seem to have overlooked this issue for decades. A notable exception is an early important contribution by Schwert (1989a and 1989b), who demonstrated how difficult it is to explain low frequency fluctuations in stock market volatility through low frequency variations in the volatility of other macroeconomic variables. Schwert, however, also showed that stock market volatility *is* countercyclical.

The main point in the present article is to revisit Schwert's contribution, and develop new empirical evidence, and theoretical explanations, in support of the view that stock market volatility is related to the business cycle, although not precisely related to the *volatility* of other macroeconomic variables. Figure 2, for example, plots return volatility against the volatility of industrial production growth, and does not reveal any statistically discernible pattern between these two volatilities. These results are in striking contrast with those available from Figure 1 in which, instead, stock market volatility exhibits a quite clear countercyclical behavior.[1] More in detail, Table 1 reveals that stock market volatility is almost 30% higher during National Bureau of Economic Research (NBER)- dated recessions than during NBER expansions.

A seemingly separate, yet very well-known, stylized fact is that risk-premia (i.e. the investors' expected return to invest in the stock market) are countercyclical (see, e.g., Fama and French (1989) and Ferson and Harvey (1991)). This fact is confirmed by the results in Table 1, which reports that expected returns are higher during NBER recessions than during NBER expansions. Table 1 also confirms one additional stylized fact uncovered by Mele (2007, 2008): the expected returns lower much less during expansions than they increase during recessions. Using post-war data, I find that compared to an average of 8.36%, the expected returns increase by nearly 19% during recessions and drop by a mere 3% during NBER expansions (see Table 1). A final stylized fact relates to the behavior of the price-dividend ratios over the business cycle. Table 1 reveals that not only are price-dividend ratios pro-cyclical. Over the last fifty years at least, price-dividend ratios movements in the

[1] The predictive regressions in Figures 1 and 2 are obtained through least absolute deviations regressions, a technique known to be more robust to the presence of outliers than ordinary least squares (see Bloomfield and Steiger (1983)).

US have also been asymmetric over the business cycle: downward changes occurring during recessions have been more severe than upward movements occurring during expansions. Table 1 suggests that price-dividend ratios fluctuate nearly two times more in recessions than in expansions.

How can we rationalize these facts? A simple possibility is that the economy is frequently hit by shocks that display the same qualitative behavior of return volatility, expected returns and price-dividend ratios. However, the empirical evidence summarized in Figure 2 suggests that this channel is unlikely. Another possibility is that the economy reacts to shocks thanks to some mechanism endogenously related to the investors' maximizing behavior, which then activates the previous phenomena.

In the first part of this essay, I develop theoretical explanations for countercyclical volatility, which rely on examples of such endogenous mechanisms. First, I explore the possibility that countercyclical volatility occurs because the investors' required return is not only countercyclical, but also *asymmetrically* related to the development of the business cycle, which happens when risk-premia increase more in bad times than they decrease in good times, as suggested by the empirical evidence summarized in Table 1. Second, I mention additional plausible explanations, which relate large price swings to the investors' process of learning about the fundamentals of the economy.

The second part of the essay provides empirical results summarizing the main cyclical properties of volatility. How can macroeconomic factors help explain the dynamics of volatility? Conversely, what is the predictive content stock market volatility brings about the development of the business cycle?

2. Understanding the empirical evidence

2.1. Fluctuating compensation for risk

In frictionless markets, the price of a long-lived security and, hence, the aggregate stock market, is simply the risk-adjusted discounted expectation of the future dividends stream. Other things being equal, this price increases as the expected return from holding the asset and hence, the risk-premium, decreases. According to this mechanism, asset prices and price-dividend ratios are pro-cyclical because risk-adjusted discount rates are countercyclical. More generally, we might think risk-premia to be one of the key elements in the engine generating the price fluctuations. Indeed, *all* the business cycle properties of asset prices must necessarily be inherited by those of the risk-premia, under rational behavior.

Let us develop intuition about why a countercyclical and asymmetric behavior of the risk-premium might lead to countercyclical volatility. Assume, first, that risk-premia are countercyclical and that they decrease less in good times than they increase in bad times, consistently with the empirical evidence discussed in Section 1. Next, suppose the economy enters a boom, in which case we expect risk-premia to decrease and asset prices to increase, on average. During the boom, the economy is hit by shocks on the fundamentals, which makes risk-premia and asset prices change. However, risk-premia and, hence, prices, do not change as they would during a recession, since we are assuming that they behave asymmetrically over the business cycle. Eventually, the boom ends and a recession begins. As the economy leads to a recession, the risk-premia increase and asset prices decrease. However, now, the shocks hitting the economy make risk-premia and, hence, prices, increase more than they decreased during the boom. Once again, the reasons for this asymmetric behavior relate to our assumption that risk-premia change asymmetrically over the business cycle.

To summarize, risk-premia are more volatile during recessions than during booms, which makes asset prices highly responsive to changes in economic conditions, in bad times. Hence, price-dividend ratios fluctuate more in bad times than in good, which translates to countercyclical return volatility. These effects are precisely those we observe, as highlighted by Table 1.

In Mele (2007), I provide a theoretical description of the previous phenomenon within a general class of models with rational expectations. A key result in that paper is that countercyclical volatility may well emerge in many equilibrium models, provided the previous asymmetry in discounting is sufficiently strong. More precisely, it is possible to show that if the asymmetry in discounting is sufficiently strong, the price-dividend ratio is an increasing and *concave* function of variables tracking the business cycle conditions. It is this concavity feature to make return volatility increase on the downside.[2]

The previous results on countercyclical volatility hold in a fairly general continuous-time framework, but their proofs are quite technical. I now offer a quantitative illustration of these results that hinges upon a simple analytical framework. I develop a tree model. This model is very simple, but it can be solved analytically and, as I show below, is able to reproduce some of the main stylized features of the actual aggregate stock market behavior.

I consider an infinite horizon economy with a representative investor who in equilibrium consumes (state by state) all the dividends promised by some asset. I assume that there exists a safe asset elastically supplied so that the safe interest rate is some constant $r > 0$. In the initial state, a dividend

[2]Under similar conditions, models with external habit formation predict counter-cyclical volatility along the same lines of arguments (see, for example, Campbell and Cochrane (1999), Menzly, Santos and Veronesi (2004), and Mele (2007)). Brunnermeier and Nagel (2007) find that US investors do not change the composition of their risky assets holdings in response to changes in wealth. The authors interpret their evidence against external habit formation. Naturally, time-varying risk-premia do not exclusively emerge in models with external habit formation. Barberis, Huang and Santos (2001) develop a theory distinct from habit formation that leads to time-varying risk-premia.

process takes a unit value (see Figure 3). In the second period, the dividend equals either $e^{-\delta}$ ($\delta > 0$) with probability p (the bad state) or e^{δ} with probability $1-p$ (the good state). In the initial state, the investor's coefficient of constant relative risk-aversion (CRRA) is $\eta > 0$. In the good (resp., the bad) state, the investor's CRRA is η_G (resp., η_B) > 0. In the third period, the investor receives the final payoffs in Figure 3, where M_S is the price of a claim to all future dividends, discounted at a CRRA η_S, with $S \in \{G, B, GB\}$ and $\eta_{GB} = \eta$ (the "hybrid" state). This model is thus one with constant expected dividend growth, but random risk-aversion.

I calibrate this model using the same US data in Table 1, and report the calibration results in Table 2. (The Appendix provides all technical details about the solution and the calibration of the model.) In spite of its overly simplifying assumptions, the model does reproduce volatility swings similar to those we observe in the data, although it might overstate the expected returns levels by some percentage points. Importantly, this calibration exercise illustrates in an exemplary manner the asymmetric feature of expected returns and risk aversion. In this simple experiment, both expected returns and risk-aversion increase much more in bad times than they decrease in good times.

2.2. Alternative stock market volatility channels

Rational explanations of stock market fluctuations must necessarily rely on some underlying state variable affecting the investors' decision environment. Two natural ways to accomplish this task are obtained through the introduction of (i) time-varying risk-premia; and (ii) time-varying expected dividend growth. The previous tree model is one simple example addressing the first extension. (More substantive examples of models predicting time-varying risk-premia are the habit formation models mentioned in footnote 2.)

Models addressing the second extension have also been produced. For example, Veronesi (1999) and Brennan and Xia (2001) have proposed models in which stock market volatility fluctuates as a result of a learning induced phenomenon. In these models, the growth rate of the economy is unknown and investors attempt to infer it from a variety of public signals. This inference process makes asset prices also depend on the investors' guesses about the dividends growth rate, and thus induces high return volatility. (In Veronesi (1999) stock market volatility is also countercyclical.) Finally, Bansal and Yaron (2004) formulated a model in which expected dividend growth is affected by some unobservable factor. This model also generates countercyclical stock market volatility. This property follows by the model's assumption that the volatilities of dividend growth and consumption are countercyclical. In contrast, in models with time-varying risk-premia (such as the previous tree model), countercyclical stock market volatility emerges without the need to impose similar features on the fundamentals of the economy. Remarkably, in models with time-varying risk-premia, countercyclical stock market volatility can be *endogenously* induced by rational fluctuations in the price-dividend ratio.

3. What to do with stock market volatility?

Both data and theory suggest that stock market volatility has a well pronounced business cycle pattern. A natural purpose at this juncture is to exploit these patterns to perform some basic forecasting exercises. I consider three in-sample exercises. First, I forecast stock market volatility from past macroeconomic data (long-run inflation, and long-run industrial production growth). Second, I forecast long-run industrial production growth from past stock market volatility. Third, I forecast the VIX index, an index of the risk adjusted expectation of future volatility, from macroeconomic data, and attempt to measure the volatility risk-premium, which is the excess amount of money a risk-averse investor is willing to pay to avoid the risk of volatility fluctuations.

3.1. Macroeconomic components of stock market volatility

Table 3 reports the results for the first forecasting exercise. Volatility is positively related to past growth. This is easy to understand. Bad times are followed by good times. Precisely, in my sample, high growth is inevitably followed by low growth. Since stock market volatility is countercyclical, high growth is followed by high stock market volatility. Stock market volatility is also related to past inflation, but in a more complex manner. Please note that once I control for past values of volatility, the results remain highly significant. Figure 4 (top panel) depicts stock market volatility and its in-sample forecasts when the regression model is fed with *past* macroeconomic data only. This fit can even be improved through the joint use of both past volatility and macroeconomic factors. Nevertheless, it is remarkable that the fit from using past macro information is more than 60% better than just using past volatility (see the R^2s in Table 3). These results are somehow in contrast with those reported in Schwert (1989). The key issue, here, remember, is that I am predicting stock market volatility within a longer time-horizon perspective.

The previous findings, while somewhat informal, suggest that relating stock market volatility to macroeconomic factors is a fertile avenue of research. The main question is, of course, how precisely stock market volatility should relate to past macroeconomic factors? Indeed, the previous linear regressions capture mere statistical relations between stock market volatility and macroeconomic factors. Yet, in the absence of arbitrage opportunities, stock market volatility is certainly related to how the price responds to shocks in the fundamentals and, hence, the macroeconomic conditions. Therefore, there should exist a no-arbitrage nexus between stock market volatility and macroeconomic factors. Corradi, Distaso and Mele (2008) pursue this topic in detail and build up a no-arbitrage model able to reproduce the previous predictability results.

3.2. Macroeconomic implications of stock market volatility

Table 4 reports results from regressing long-run growth on to macroeconomic variables and return volatility (only R^2s are reported). The volatility concept I use is purely related to volatility induced by price-dividend fluctuations (i.e. it is *not* related to dividend growth volatility). I find that the predictive power of traditional macroeconomic variables is considerably enhanced (almost doubled) with the inclusion of this new volatility concept and the price-dividend ratio. According to Figure 4 (bottom panel), stock market volatility does help predicting the business cycle. Fornari and Mele (2006) contain many details on the forecasting performance of a new block, including stock market volatility and the slope of the yield curve. They show this block is quite successful and outperforms traditional models based on financial variables, both in sample and out of sample.

3.3. Risk-adjusted volatility

3.3.1. Volatility trading

An important financial innovation for volatility trading was the introduction of the "variance swaps", which are contracts allowing to trade future realized variance against a fixed swap rate. Passive funds managers find this contract useful, as in times of high volatility, tracking errors widen and, then, funds performance deteriorates. Many hedge funds find this type of contracts attractive as well, as these funds invest in "relative value" strategies, attempting to profit from temporary price discrepancies. In times of high volatility, price discrepancies typically widen, and volatility contracts help these institutions to hedge against these events.

All in all, the payoff guaranteed to the buyer of a swap equals the difference between the realized volatility over the life of the contract and a fixed swap rate. Entering this contract at time of origination does not cost. Therefore, the fixed swap rate is equal to the expectation of the future realized volatility under the risk-neutral probability. In September 2003, the CBOE started to calculate the VIX index in a way that makes this index equal to such a risk-neutral expectation. The strength of this index is that although it deals with risk-neutral expectations, it is nonparametric - it does not rely on any model of stochastic volatility. Rather, it is based on a basket of all the available options, following the seminal work by Demeterfi et al. (1999), Bakshi and Madan (2000), Britten-Jones and Neuberger (2000), and Carr and Madan (2001).

3.3.2. Business cycle determinants of volatility trading

Figure 5 (top panel) depicts the VIX index, along with predictions obtained through a parametric model. The predicting model is based on the regression of the VIX index on the same macroeconomic variables considered in the previous sections: inflation and growth. Table 5 reports the estimation results, which reveal how important the contribution of macroeconomic factors is to explain the dynamics of the VIX.

Finally, Figure 5 (bottom panel) depicts the volatility risk-premium, defined as the difference between the expectation of future volatility under the risk-neutral and the physical probability. I estimated the risk-neutral expectation as the predicting part of the linear regression on macroeconomic factors (the dotted line in Figure 5, top panel). I estimated expected volatility as the predicting part of an AR(1) model fitted to the volatility depicted in Figure 4 (top panel). As we see, volatility risk-premia are indeed strongly countercyclical. Once again, the results in these picture are very suggestive, but they do represent mere statistical relations. The model considered by Corradi, Distaso and Mele (2008) has the strength to make these statistical relations emerge as a result of a fully articulated no-arbitrage model.

4. Conclusions

Stock market volatility is higher in bad times than in good times. Explaining this basic fact is challenging. Indeed, economists know very well how to model risk-premia and how these premia should relate to the business cycle. We feel more embarrassed when we come to explain volatility. The ambition in this short essay is to explain that countercyclical volatility can be made consistent with the prediction of the neoclassical model of asset pricing - in which asset prices are (risk-adjusted) expectations of future dividends. One condition activating countercyclical volatility is very simple: risk-premia must swing sharply as the economy moves away from good states, just as the data seem to suggest. The focus in this paper is about stock market volatility *fluctuations*. Accordingly, I did not discuss whether the *average levels* of stock market volatility and risk-premia are consistent with plausible levels of investors' risk-aversion, although I note that these issues could be as controversial as the traditional topics at the intersection between financial economics and macroeconomics (see, e.g., Campbell (2003) and Mehra and Prescott (2003)). However, as I demonstrated with a basic tree, the neoclassical model seems promising in explaining how volatility switches across states.

The final contribution of this essay is to investigate whether these theoretical insights have some additional empirical content. I explored three empirical issues, which are currently being investigated in the literature. First, I demonstrated that stock market volatility can be forecast through macroeconomic variables. Second, I showed that in turn, stock market volatility contains relevant information related to the development of the business cycle. Third, I demonstrated that volatility trading is related to the business cycle and that volatility risk-premia are strongly countercyclical.

Appendix: Calibration of the tree

Solution and calibration of the model. The initial step of the calibration reported in Table 2 involves estimating the two parameters p and δ of the dividend process. Let G be the dividend gross growth rate, computed at a yearly frequency. I calibrate p and δ by a perfect matching of the model's expected dividend growth, $\mu_D \equiv E(G) = pe^{-\delta} + (1-p)e^{\delta}$, and the model's dividend variance, $\sigma_D^2 \equiv var(G) = \left(e^{\delta} - e^{-\delta}\right)^2 p(1-p)$, to their sample counterparts $\hat{\mu}_D = 1.0594$ and $\hat{\sigma}_D = 0.0602$ obtained on US aggregate dividend data. The result is $(p, \delta) = (0.158, 0.082)$. Given these calibrated values of (p, δ), I fix $r = 1.0\%$, and proceed to calibrate the probabilities q, q_B and q_G.

To calibrate (q, q_B, q_G), I need an explicit expression for all the payoffs at each node. By standard risk-neutral evaluation, I obtain a closed form solution for the price of the claim M_S, as follows. For each state $S \in \{G, B, GB\}$, M_S is solution to,

$$\frac{M_S}{D_S} = e^{-r}\mathbb{E}_S\left(\frac{D_S'}{D_S} + \frac{M_S'}{D_S'}\frac{D_S'}{D_S}\right), \tag{A.1}$$

where $\mathbb{E}_S(\cdot)$ is the expectation taken under the risk-neutral probability q_S in state S, $S \in \{G, B, GB\}$, and $q_{GB} = q$, $D_G = e^{2\delta}$, $D_B = e^{-2\delta}$, $D_{GB} = 1$, and D_S' and M_S' are the dividend and the price of the claim as of the next period. Since risk-aversion is constant from the third period on, the price-dividend ratio is constant as well, from the third period on, which implies that $\frac{M_S}{D_S} = \frac{M_S'}{D_S'}$. By using the equality $\frac{M_S}{D_S} = \frac{M_S'}{D_S'}$ in Eq. (A.1), and solving for M_S, yields,

$$M_S = D_S \frac{q_S e^{-\delta} + (1-q_S)e^{\delta}}{e^r - \left[q_S e^{-\delta} + (1-q_S)e^{\delta}\right]}, \quad S \in \{G, B, GB\}. \tag{A.2}$$

I calibrate $(q_G, q_B, q_{GB} = q)$ to make the "hybrid" price-dividend (P/D henceforth) ratio M_{GB}, the "good" P/D ratio $\frac{M_G}{e^{2\delta}}$ and the "bad" P/D ratio $\frac{M_B}{e^{-2\delta}}$ in Eq. (A.2) perfectly match the average P/D ratio, the average P/D ratio during NBER expansion periods, and the average P/D ratio during NBER recession periods (i.e. 31.99, 33.21 and 26.20, from Table 1). Given $(p, \delta, r, q, q_S, q_G)$, I compute the P/D ratios in states G and B. For example, the price of the asset in state B is, $P_B = e^{-r}[q_B\left(e^{-2\delta} + M_B\right) + (1-q_B)(1 + M_{GB})]$. Given P_B, I compute the log-return in the bad state as $\log(\frac{\tilde{\Pi}}{P_B})$, where either $\tilde{\Pi} = e^{-2\delta} + M_B$ with probability p, or $\tilde{\Pi} = 1 + M_{GB}$ with probability $1-p$. Then, I compute the return volatility in state B. The P/D ratios, the expected log-return and return volatility in state G are computed similarly. (Please notice that volatilities under p and under $\{q_S\}_{S\in\{G,B,GB\}}$ are not the same.)

Next, I recover the risk-aversion parameter η_S in the three states $S \in \{G, B, GB\}$ implied by the previously calibrated probabilities q, q_G and $q = q_{GB}$. As I shall show below, the relevant formula to use is,

$$\frac{q_S}{p} = \frac{e^{\eta_S\delta}}{pe^{\eta_S\delta} + (1-p)e^{-\eta_S\delta}}, \quad S \in \{G, B, GB\}. \tag{A.3}$$

The values for the "implied" risk-aversion parameter in Table 2 are obtained by inverting Eq. (A.3) for η_S, given the calibrated values of (p, δ, q_S, q_G).

Finally, I compute the risk-adjusted discount rate as $r + \hat{\sigma}_D\lambda_S$, where λ_S is the Sharpe ratio, which I shall show below to equal,

$$\lambda_S = \frac{q_S - p}{\sqrt{p(1-p)}}, \quad S \in \{G, B, GB\}. \tag{A.4}$$

Proof of Eq. (A.3). I only provide the derivation of the risk-neutral probability q_B, since the proofs for the expressions of the risk-neutral probabilities q_G and $q = q_{GB}$ are nearly identical. In equilibrium, the Euler equation for the stock price at the "bad" node is,

$$P_B = \beta E\left[\frac{u'_B(\tilde{D}_S)}{u'_B(e^{-\delta})}\left(\tilde{D}_S + M_S\right)\right] = \beta E\left[\tilde{G}_S^{-\eta_B}\left(\tilde{D}_S + M_S\right)\right], \quad S \in \{B, GB\}, \tag{A.5}$$

where: (i) β is the discount rate; (ii) the utility function for consumption C is state dependent and equal to, $u_B(C) = C^{1-\eta_B}/(1-\eta_B)$; (iii) $E(\cdot)$ is the expectation taken under the probability p; and (iv) the dividend $\tilde{D}_S$ and the gross dividend growth rate $\tilde{G}_S$ are either $\tilde{D}_B = e^{-2\delta}$ and $\tilde{G}_B = \frac{e^{-2\delta}}{e^{-\delta}} = e^{-\delta}$ with probability p, or $\tilde{D}_{GB} = 1$ and $\tilde{G}_{GB} = \frac{1}{e^{-\delta}} = e^{\delta}$ with probability $1-p$.

The model I set up assumes that the asset is elastically supplied or, equivalently, that there exists a storage technology with a fixed rate of return equal to $r = 1\%$. Let us derive the agent's private evaluation of this asset. The Euler equation for the safe asset is,

$$e^{-r_B} = \beta E[\tilde{G}_S^{-\eta_B}] = \beta \sum_{S\in\{B,GB\}} p_S \tilde{G}_S^{-\eta_B}, \tag{A.6}$$

where the safe interest rate, r_B, is state dependent, $p_B = p$ and $p_{GB} = 1-p$. Therefore,

$$q_B = \beta e^{r_B} p \tilde{G}_B^{-\eta_B} \quad ; \quad 1 - q_B = \beta e^{r_B}(1-p)\,\tilde{G}_{GB}^{-\eta_B} \tag{A.7}$$

is a probability distribution. In fact, by plugging q_B and $1-q_B$ into Eq. (A.5), one sees that it is the risk-neutral probability distribution. To obtain Eq. (A.3), note that by Eq. (A.6), $\beta e^{r_B} = 1/E[\tilde{G}_S^{-\eta_B}]$, which replaced into Eq. (A.7) yields,

$$\frac{q_B}{p} = \frac{\tilde{G}_B^{-\eta_B}}{E[\tilde{G}_S^{-\eta_B}]}.$$

Eq. (A.3) follows by the definition of $\tilde{G}_S$ given above.

Proof of Eq. (A.4). Let e^{μ} the gross expected return of the risky asset. The asset return can take two values: e^{R_ℓ} with probability p, and e^{R_h} with probability $1-p$, and $R_h > R_\ell$. Therefore, for each state, we have that:

$$e^{\mu} = pe^{R_\ell} + (1-p)\,e^{R_h} \quad ; \quad e^{r} = qe^{R_\ell} + (1-q)\,e^{R_h}, \tag{A.8}$$

where we have omitted the dependence on the state S to alleviate the presentation. The standard deviation of the asset return is $\text{Std}_R = \left(e^{R_h} - e^{R_\ell}\right)\sqrt{p(1-p)}$. The Sharpe ratio is defined as

$$\lambda = \frac{e^{\mu} - e^{r}}{\text{Std}_R}.$$

By substracting the two equations in (A.8),

$$q = p + \frac{e^{\mu} - e^{r}}{\left(e^{R_h} - e^{R_\ell}\right)\sqrt{p(1-p)}}\sqrt{p(1-p)} = p - \lambda\sqrt{p(1-p)},$$

from which Eq. (A.4) follows immediately. Note, also, that in terms of this definition of the Sharpe ratio, the risk-neutral expectation of the dividend growth is, $\mathbb{E}(G) = E(G) - \lambda\sigma_D$.

References

Andersen, T. G., T. Bollerslev and F. X. Diebold, 2002. "Parametric and Nonparametric Volatility Measurement." Forthcoming in Aït-Sahalia, Y. and L. P. Hansen (Eds.): *Handbook of Financial Econometrics.*

Bakshi, G. and D. Madan, 2000. "Spanning and Derivative Security Evaluation." *Journal of Financial Economics* 55, 205-238.

Bansal, R. and A. Yaron, 2004. "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles." *Journal of Finance* 59, 1481-1509.

Barberis, N., M. Huang and T. Santos, 2001. "Prospect Theory and Asset Prices." *Quarterly Journal of Economics* 116, 1-53.

Bloomfield, P. and Steiger, W. L., 1983. *Least Absolute Deviations.* Boston: Birkhäuser.

Brennan, M. J. and Y. Xia, 2001. "Stock Price Volatility and Equity Premium." *Journal of Monetary Economics* 47, 249-283.

Britten-Jones, M. and A. Neuberger, 2000. "Option Prices, Implied Price Processes and Stochastic Volatility." *Journal of Finance* 55, 839-866.

Brunnermeier, M. K. and S. Nagel, 2007. "Do Wealth Fluctuations Generate Time-Varying Risk Aversion? Micro-Evidence on Individuals' Asset Allocation." Forthcoming in *American Economic Review.*

Campbell, J. Y., 2003. "Consumption-Based Asset Pricing." In Constantinides, G.M., M. Harris and R. M. Stulz (Editors): *Handbook of the Economics of Finance* (Volume 1B, Chapter 13), 803-887.

Campbell, J. Y., and J. H. Cochrane, 1999. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *Journal of Political Economy* 107, 205-251.

Carr, P. and D. Madan, 2001. "Optimal Positioning in Derivative Securities." *Quantitative Finance* 1, 19-37.

Corradi, V., W. Distaso and A. Mele, 2008. "Macroeconomic Determinants of Stock Market Volatility and Volatility Risk-Premia." Working paper University of Warwick, Imperial College, and London School of Economics.

Demeterfi, K., E. Derman, M. Kamal and J. Zou, 1999. "A Guide to Volatility and Variance Swaps." *Journal of Derivatives* 6, 9-32.

Fama, E. F. and K. R. French, 1989. "Business Conditions and Expected Returns on Stock and Bonds." *Journal of Financial Economics* 25, 23-49.

Ferson, W. E. and C. R. Harvey, 1991. "The Variation of Economic Risk Premiums." *Journal of Political Economy* 99, 385-415.

Fornari, F. and A. Mele, 2006. "Financial Volatility and Real Economic Activity." Working paper European Central Bank and London School of Economics.

Lewis, A. L., 2000. *Option Valuation Under Stochastic Volatility.* Finance Press, Newport Beach, CA.

Mehra, R. and E. C. Prescott, 2003. "The Equity Premium in Retrospect." In Constantinides, G.M., M. Harris and R. M. Stulz (Editors): *Handbook of the Economics of Finance* (Volume 1B, chapter 14), 889-938.

Mele, A., 2007. "Asymmetric Stock Market Volatility and the Cyclical Behavior of Expected Returns." *Journal of Financial Economics* 86, 446-478.

Mele, A., 2008. "Aggregate Stock Market Risk-Premia and Real Economic Activity." Working paper London School of Economics.

Menzly, L., T. Santos and P. Veronesi, 2004, "Understanding Predictability." *Journal of Political Economy* 111, 1, 1-47.

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." *Carnegie-Rochester Conference Series on Public Policy* 31, 83-125.

Veronesi, P., 1999, "Stock Market Overreaction to Bad News in Good Times: A Rational Expectations Equilibrium Model." *Review of Financial Studies* 12, 975-1007.

Tables

Table 1 – Business cycle properties of P/D ratios and returns

	total		NBER expansions		NBER recessions	
	average	std dev	average	std dev	average	std dev
P/D ratio	31.99	15.88	33.21	15.79	26.20	14.89
$\ln \frac{\text{P/D}_{t+1}}{\text{P/D}_t}$	2.01	12.13	3.95	10.81	−7.28	16.79
one year returns	8.59	15.86	12.41	13.04	−9.45	15.49
real risk-free rate	1.02	2.48	1.03	2.43	0.97	2.69
excess return volatility	11.34	3.89	10.80	3.59	13.91	4.15
expected returns	8.36	3.49	8.09	3.29	9.62	4.10

P/D is the S&P Comp. price-dividend ratio. One year returns as of time t are computed as $\sum_{i=1}^{12} \tilde{R}_{t-i}$, where $\tilde{R}_t$ are the (log-)returns on the S&P Comp. return deflated by the CPI, in excess of the real (one month) risk-free rate. Expected returns are computed through the Fama & French (1989) predictive regressions of $\tilde{R}_t$ on to default-premium, term-premium and return volatility (defined below). Volatility is the excess return volatility, and is computed as $\text{Vol}_t \equiv (|\text{exc}_t| + \cdots + |\text{exc}_{t-11}|)/\sqrt{12}$, where exc_t is the S&P Comp. demeaned return in excess of the riskless asset, as of month t. Data are sampled monthly and cover the period from January 1948 through December 2002. With the exception of the P/D ratio levels, all figures are annualized percent.

Table 2 – Infinite horizon model

	Data		
	expansions	average	recessions
P/D ratio	33.21	31.99	26.20
excess return volatility	10.80	11.34	13.91

	Model calibration		
	good state	average	bad state
P/D ratio	32.50	31.81	28.15
excess return volatility	7.29	8.20	13.03
risk-adjusted rate	8.95	9.07	9.71
expected returns	10.16	11.46	18.42
implied risk-aversion	13.69	13.89	14.96

This table reports calibration results for the infinite horizon tree model in Figure 3. The expected returns and excess return volatility predicted by the model are computed using log-returns. The risk-adjusted rate is computed as $r + \hat{\sigma}_D \lambda_S$, where: r is the continuously compounded riskless rate; $\hat{\sigma}_D$ is the dividend volatility; λ_S is the Sharpe ratio on gross returns in state S, computed as $\lambda_S \equiv (q_S - p) \div \sqrt{p(1-p)}$ for $S = G$ (the good state) and $S = B$ (the bad state); p is the probability of the bad state; and q_S is the state dependent risk-adjusted probability of a bad state (for $S \in \{G, B\}$). Implied risk-aversion is the coefficient of relative risk aversion η_S in the good state ($S = G$) and in the bad state ($S = B$), implied by the calibrated model. The figures in the "average" column are the averages of the corresponding values in the good and bad states taken under the probability $p = 0.158$.

Table 3 – Forecasting stock market volatility with economic activity

Past				Future	
Const.	6.92	7.76	2.48	Const.	8.28
Growth_{t-12}	–	0.29*	1.67	Growth_{t+12}	0.21*
Growth_{t-24}	–	0.74	1.09	Growth_{t+24}	1.62
Growth_{t-36}	–	2.17	2.44	Growth_{t+36}	−0.02*
Growth_{t-48}	–	1.77	1.91	Growth_{t+48}	0.12*
Infl_{t-12}	–	10.44	8.05	Infl_{t+12}	3.55
Infl_{t-24}	–	−5.96	−5.49	Infl_{t+24}	−0.81*
Infl_{t-36}	–	−1.42*	−0.97	Infl_{t+36}	−0.54*
Infl_{t-48}	–	3.73	3.31	Infl_{t+48}	4.33
Vol_{t-12}	0.43	–	0.37	R^2	12.79
Vol_{t-24}	−0.17	–	−0.09		
Vol_{t-36}	0.02*	–	0.09		
Vol_{t-48}	0.12	–	0.09		
R^2	16.38	26.01	34.52		

The first part of this table ("Past") reports ordinary least square coefficient estimates in linear regression of volatility on to, *past* long-run industrial production growth (defined in Figure 1), *past* long-run inflation (defined similarly as in Figure 1), and *past* long-run volatility. Growth_{t-12} is the long-run industrial production growth at time $t-12$, etc. Time units are months. The second part of the table ("Future") is similar, but it contains coefficient estimates in linear regression of volatility on to *future* long-run industrial production growth and *future* long-run inflation. Starred figures are not statistically distinguishable from zero at the 95% level. R^2 is the percentage, adjusted R^2.

Table 4 – Forecasting economic activity with stock market volatility

Predictors	R^2
(i) P/D Volatility	10.81
(ii) P/D ratio	15.57
(iii) P/D Volatility, P/D ratio	20.98
(iv) Growth, Inflation	21.20
(v) Growth, Inflation, P/D volatility	34.29
(vi) Growth, Inflation, P/D volatility, P/D ratio	41.76

This table reports the R^2 (adjusted, in percentage) from six linear regressions of 6-months moving average industrial production growth on to the listed set of predicting variables. Inflation is also 6-months moving average inflation. The regressors lags are 6-months, and 1, 2 and 3 years. P/D volatility is defined as a 12 months moving average of $abs(\log(\frac{1+\mathrm{P/D}_{t+1}}{\mathrm{P/D}_t}))$, where $abs(\cdot)$ denotes the absolute value, and P/D is the price-dividend ratio.

Table 5 – Forecasting the VIX index with economic activity

	Past		
Const.	2.60*	30.15	3.03*
Growth$_{t-1}$	–	-5.12	0.51*
Growth$_{t-12}$	–	-3.69*	-0.35*
Growth$_{t-24}$	–	4.91	3.69
Growth$_{t-36}$	–	11.19	4.33
Infl$_{t-1}$	–	-26.96	-9.14*
Infl$_{t-12}$	–	-22.62	-1.89*
Infl$_{t-24}$	–	-1.59*	5.85*
Infl$_{t-36}$	–	-6.02*	-2.56*
VIX$_{t-1}$	0.72	–	0.55
VIX$_{t-12}$	0.18	–	0.14*
VIX$_{t-24}$	-0.06*	–	-0.01*
VIX$_{t-36}$	0.02*	–	0.12*
R^2	66.87	54.12	71.03

	Future
Const.	25.53
Growth$_{t+1}$	25.53
Growth$_{t+12}$	-5.58
Growth$_{t+24}$	-8.34
Growth$_{t+36}$	-9.67
Infl$_{t+1}$	1.04*
Infl$_{t+12}$	-24.11
Infl$_{t+24}$	9.32*
Infl$_{t+36}$	20.71
R^2	55.04

The first part of this table ("Past") reports ordinary least square coefficient estimates in linear regression of the VIX index on to, *past* long-run industrial production growth (defined in Figure 1), *past* long-run inflation (defined similarly as in Figure 1), and *past* long-run volatility. Growth$_{t-12}$ is the long-run industrial production growth at time $t-12$, etc. Time units are months. The second part of the table ("Future") is similar, but it contains coefficient estimates in linear regression of the VIX index on to *future* long-run industrial production growth and *future* long-run inflation. Starred figures are not statistically distinguishable from zero at the 95% level. R^2 is the percentage, adjusted R^2.

Figures

Figure 1 – Return volatility and business cycle conditions

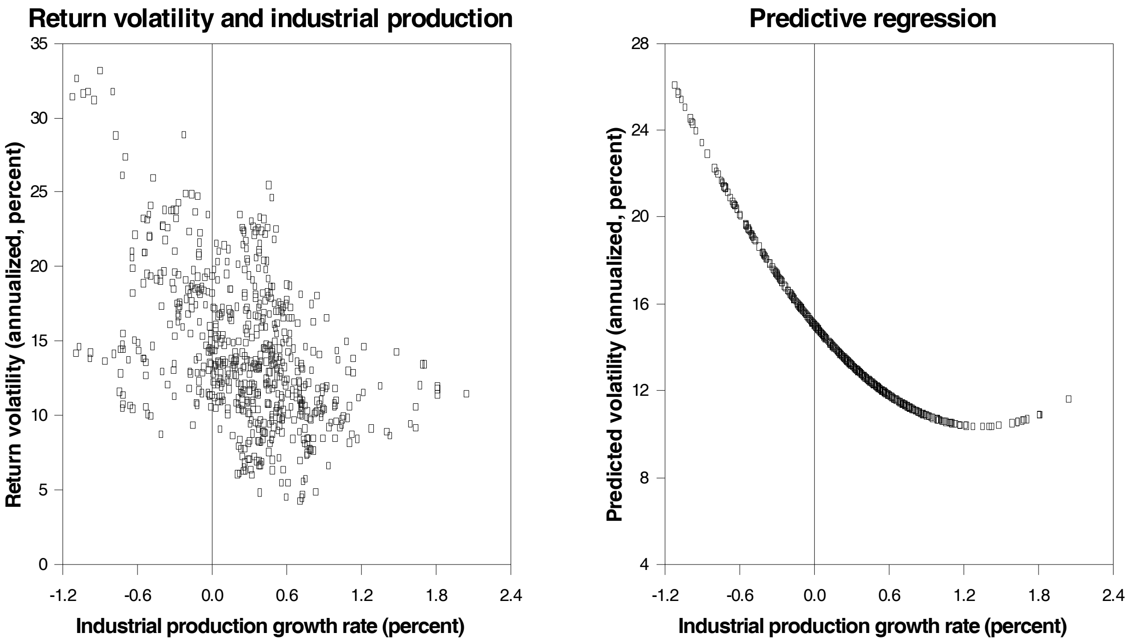

The left panel plots a measure of stock return volatility against a measure of long-run industrial production growth. Return volatility is computed as $\mathrm{Vol}_t \equiv \left(|\mathrm{exc}_t| + \cdots + |\mathrm{exc}_{t-11}|\right)/\sqrt{12}$, where exc_t is the S&P Comp. demeaned return in excess of the riskless asset, as of month t. Long-run industrial production growth is computed as $\mathrm{IP}_t \equiv (\mathrm{Ind}_t + \cdots + \mathrm{Ind}_{t-11})/12$, where Ind_t is the real, seasonally adjusted industrial production growth rate as of month t. The right panel depicts the prediction of the static least absolute deviations regression: $\mathrm{Vol} = \underset{(0.16)}{12.01} - \underset{(0.33)}{5.57} \cdot \mathrm{IP} + \underset{(0.35)}{2.06} \cdot \mathrm{IP}^2 + w$, where w is a residual term, and standard errors are in parenthesis. The data span the period from January 1948 to December 2002.

Figure 2 – Return volatility and industrial production volatility

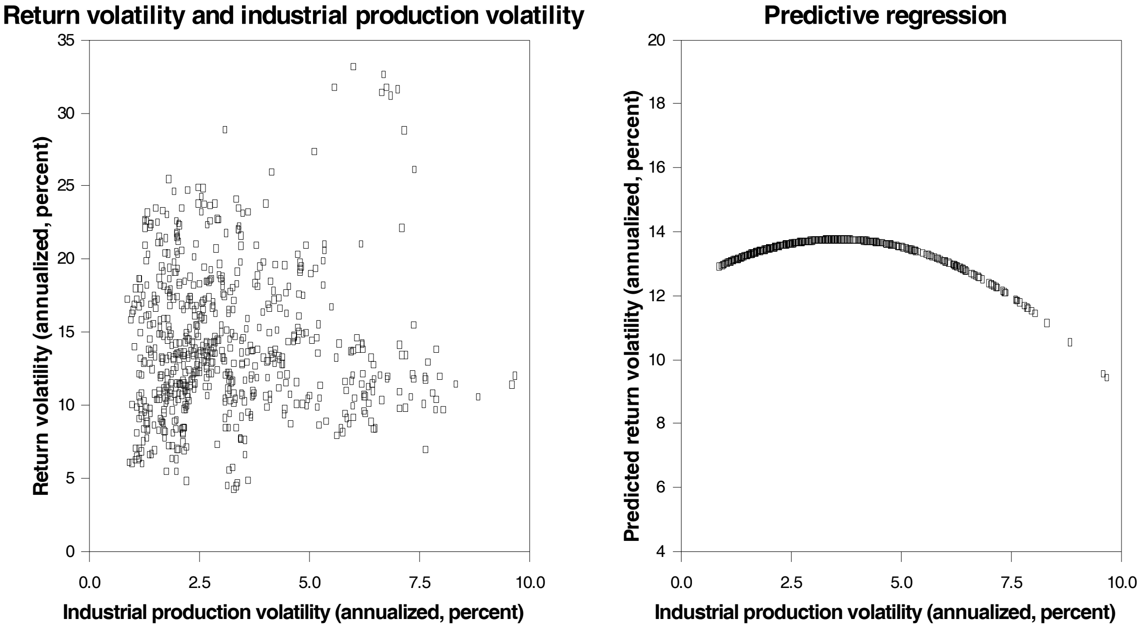

The left panel plots a measure of stock return volatility against a measure of industrial production volatility. Return volatility is computed as $\text{Vol}_t \equiv (\,|\text{exc}_t| + \cdots + |\text{exc}_{t-11}|)/\sqrt{12}$, where exc_t is the demeaned return in excess of the S&P Comp. demeaned return in excess of the riskless asset, as of month t. Industrial production volatility is computed as $\text{Vol}_{\text{G},t} \equiv (\,|G_t| + \cdots + |G_{t-11}|)/\sqrt{12}$, where G_t is the real, seasonally adjusted industrial production growth rate as of month t. The right panel of the picture depicts the prediction of the static least absolute deviations regression: $\text{Vol} = \underset{(0.83)}{12.28} - \underset{(0.47)}{0.83} \cdot \text{Vol}_{\text{G}} - \underset{(0.05)}{0.12} \cdot \text{Vol}_{\text{G}}^2 + w$, where w is a residual term, and standard errors are in parenthesis. The data span the period from January 1948 to December 2002.

Figure 3 – A tree model of random risk-aversion and countercyclical volatility

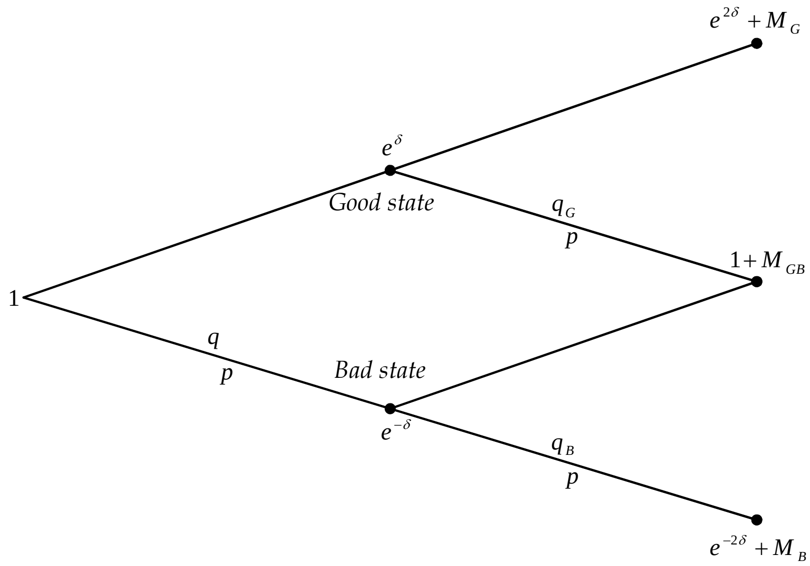

The dividend process takes a unit value at the initial node. With probability p, the dividend then decreases to $e^{-\delta}$ in the bad state. The corresponding risk-neutral probability is denoted as q. The risk-neutral probability of further dividends movements differs according to whether the economy is in the good or bad state (i.e. q_G or q_B). At the end of the tree, the investor receives the dividends plus the right to the stream of all future dividends. In the upper node, this right is worth M_G (the evaluation obtained through the risk-neutral probability q_G). In the central node it is worth M_{GB} (the evaluation obtained through the risk-neutral probability q). In the lower node it is worth M_B (the evaluation obtained through the risk-neutral probability q_B). The safe interest rate is constant in this model.

Figure 4 – Forecasts

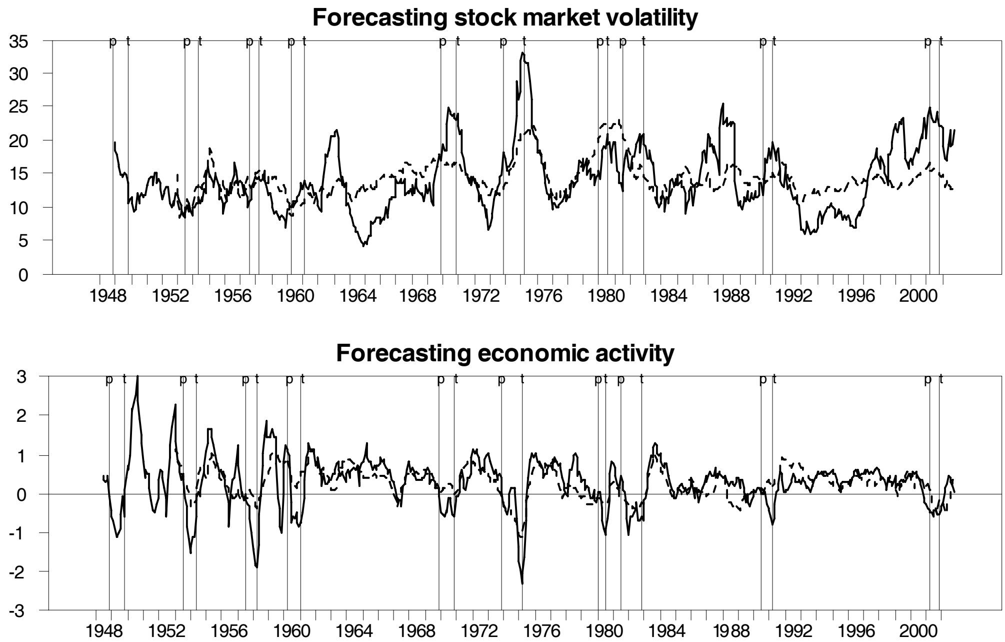

The top panel depicts stock market volatility (solid line) and stock market volatility forecasts obtained through the sole use of the macroeconomic indicators in Table 3 (dashed line). Stock market volatility is defined as in Figure 1. The bottom panel depicts 6-months moving average industrial production growth (solid line) and its forecasts based on the 6th regression in Table 4 (dashed line). Vertical lines marked with a 'p' (peak) track the beginning of NBER-dated recessions, and vertical lines marked with a 't' (trough) indicate the end of NBER-dated recessions.

Figure 5 – Forecasting the VIX Index, and the volatility risk-premium

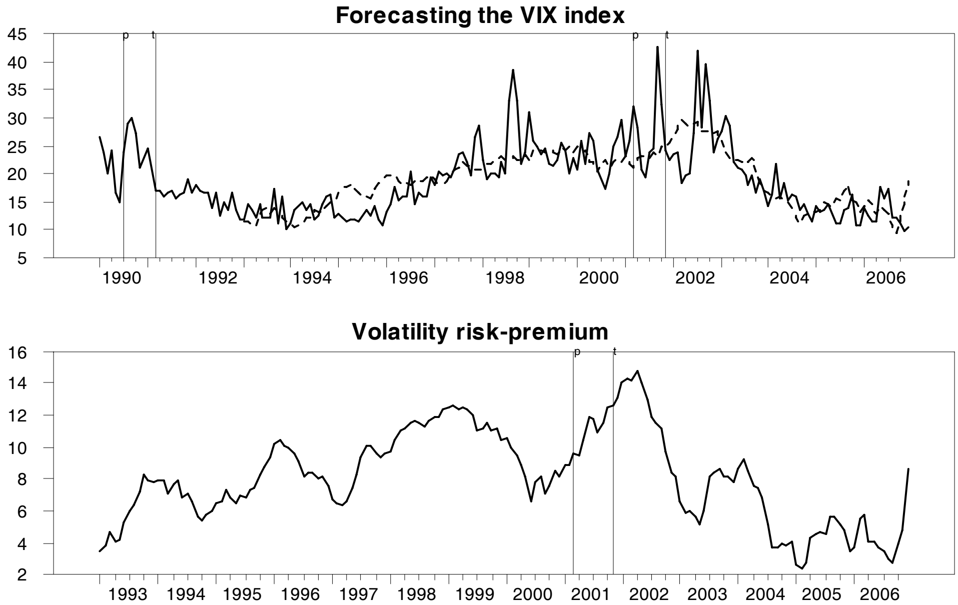

The top panel depicts the VIX index (solid line) and the VIX forecasts obtained through the sole use of the macroeconomic indicators in Table 5 (dashed line). The bottom panel plots the volatility risk premium, defined as the difference between the one-month ahead volatility forecast calculated under the risk-neutral probability and the one-month ahead volatility forecast calculated under the physical probability. Vertical lines marked with a 'p' (peak) track the beginning of NBER-dated recessions, and vertical lines marked with a 't' (trough) indicate the end of NBER-dated recessions.