European Economic Review 70 (2014) 88–107

Contents lists available at ScienceDirect

European Economic Review

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

Risk aversion, investor information and stock market volatility ☆

Kevin J. Lansing [a,*], Stephen F. LeRoy [a,b,1]

[a] Research Department, Federal Reserve Bank of San Francisco, P.O. Box 7702, San Francisco, CA 94120-7702, United States
[b] Department of Economics, University of California, Santa Barbara, CA 93106, United States

ARTICLE INFO

Article history:
Received 14 May 2013
Accepted 17 March 2014
Available online 5 April 2014

JEL classification:
E44
G12

Keywords:
Asset pricing
Excess volatility
Variance bounds
Risk aversion
Imperfect information

ABSTRACT

This paper employs a standard asset pricing model to derive theoretical volatility measures in a setting that allows for varying degrees of investor information about the dividend process. We show that the volatility of the price–dividend ratio increases monotonically with investor information but the relationship between investor information and equity return volatility (or equity premium volatility) can be non-monotonic, depending on risk aversion and other parameter values. Under some plausible calibrations and information assumptions, we show that the model can match the standard deviations of equity market variables in long-run U.S. data. In the absence of concrete knowledge about investors' information, it becomes more difficult to conclude that observed volatility in the data is excessive.

1. Introduction

Numerous empirical studies starting with Shiller (1981) and LeRoy and Porter (1981) concluded that U.S. stock market volatility appeared excessive when compared to the present-value of ex post realized dividends discounted at a constant rate, implying risk-neutral investors. A number of econometric problems with the empirical studies were later raised (e.g., Kleidon, 1986; Marsh and Merton, 1986), but it turned out that correcting these problems did not eliminate the appearance of excess volatility.[2]

Other studies around this time (e.g., Grossman and Shiller, 1981; LeRoy and LaCivita, 1981) recognized that allowing for risk aversion when discounting the stream of ex post realized dividends could increase volatility relative to the risk-neutral case. However, the hypothetical stock price series computed in this way was still only linked to a single information assumption, i.e., perfect foresight on the part of investors about the path of future dividends.

In this paper, we employ a standard Lucas-type asset pricing model with power utility and exponentially growing dividends to derive theoretical volatility measures in a setting that allows for varying degrees of investor information about the dividend process. We examine four different information sets labeled G_t, H_t, J_t, and I_t^* that contain progressively increasing amounts of information, i.e., $G_t \subseteq H_t \subseteq J_t \subseteq I_t^*$. Under set G_t, the investor can observe current and past dividend

☆ An earlier version of this paper was titled "Risk Aversion and Stock Price Volatility."
* Corresponding author. Tel.: +1 415 974 2393; fax: +1 415 977 4031.
E-mail addresses: kevin.j.lansing@sf.frb.org (K.J. Lansing), sleroy@econ.ucsb.edu (S.F. LeRoy).
URLS: http://www.frbsf.org/economics/economists/klansing.html (K.J. Lansing), http://www.econ.ucsb.edu/~sleroy/webpage/ (S.F. LeRoy).
[1] Tel.: +1 805 689 2344.
[2] For summaries of this extensive literature, see West (1988a), Gilles and LeRoy (1991), Shiller (2003), and LeRoy (2010).

http://dx.doi.org/10.1016/j.euroecorev.2014.03.009
0014-2921/

realizations but observations of trend dividend growth are contaminated with noise. This imperfect information setup is similar to one considered by Veronesi (2000).[3] Set H_t provides more information than set G_t by allowing investors to directly observe trend growth and thereby identify the noise shocks. Set J_t goes a step further by allowing investors to have one-period foresight about dividends and the trend growth rate. This setup captures the possibility that investors may have some auxiliary information that allows them to accurately forecast dividends and trend growth over the near-term. Information set J_t connects to recent research on business cycles that focuses on "news shocks" as an important quantitative source of economic fluctuations. In these models, news shocks provide agents with auxiliary information about future technology innovations.[4] Finally, set I_t^* provides the maximum amount of investor information, corresponding to perfect knowledge about the entire stream of past and future dividends and trend growth rates. While this information assumption is obviously extreme, it provides a useful benchmark and helps connect our results to the earlier literature on stock market volatility mentioned above.

We demonstrate that the assumed degree of investor information can have significant qualitative and quantitative impacts on the volatility of equity market variables in the model. The volatility of the price–dividend ratio increases monotonically with investor information but the relationship between investor information and equity return volatility can be non-monotonic, depending on risk aversion and other parameter values. Put differently, providing investors with more information about the dividend process can either increase or decrease the volatility of the equity return. There also can be a non-monotonic relationship between investor information and the volatility of the excess return on equity, i.e., the equity premium.

The intuition for the complex relationship between investor information and return volatility is linked to the discounting mechanism. Two crucial elements are the persistence of trend dividend growth and the investor's discount factor (which depends on the coefficient of relative risk aversion). Both elements affect the degree to which future dividend innovations influence the perfect foresight price via discounting from the future to the present. When dividends are sufficiently persistent and the investor's discount factor is sufficiently close to unity, the discounting weights applied to successive future dividend innovations decay gradually. Since log returns are nearly the same as log price-changes, computation of the log return under information set I_t^* tends to "difference out" the future dividend innovations, thus shrinking the magnitude of the perfect foresight return variance relative to the other information sets. In contrast, when dividend growth is less persistent and/or the investor's discount factor is much less than unity, the discounting weights applied to successive future dividend innovations decay rapidly. Consequently, these terms do not tend to difference out which serves to magnify the perfect foresight return variance relative to the other information sets. Similar logic applies when comparing return volatility under information set J_t (one-period foresight) to return volatility under information sets G_t or H_t.

The log return variance in our model is the analog to the arithmetic price-change variance examined by West (1988b) and Engel (2005) in risk-neutral settings with arithmetically growing dividends. They show that the arithmetic price-change variance is a monotonically decreasing function of investors' information about future dividends. In contrast, we show that when investors are risk averse, the analogs to the West–Engel results do not go through; log return variance (or log price-change variance) is not a monotonic decreasing function of investors' information about future dividends.[5] Our results have implications for the behavior of other asset prices, such as exchange rates. For example, Engel (2014, p. 11) states "…the variance of changes in the asset price falls with more information…[N]ews can account for a high variance in the real exchange rate, but not for a high variance in the change in the real exchange rate." Our results demonstrate that the variance of log returns (or log price-changes) can rise with more information, thereby allowing new shocks to help account for the high variance of exchange rate changes or other asset price changes.

As part of our quantitative analysis, we compare model-predicted volatilities to the corresponding values in long-run U.S. stock market data. Using plausible calibrations for the noisy dividend process and the coefficient of relative risk aversion, we show that some specifications of the model can match the standard deviations of the log price dividend ratio, the log equity return, and the log excess return on equity in the data. For the baseline calibration, model-predicted volatility for the log price–dividend ratio can match the data only when investors are endowed with at least some knowledge about future dividends, i.e., information sets J_t or I_t^*. The perfect foresight case requires a coefficient of relative risk aversion around 4 to match the data volatility. However, in Section 5 of the paper, we show that the model under information set G_t (least information) can match the data volatility with a risk aversion coefficient around 5 if we allow for a highly persistent trend growth process combined with a more volatile noise shock (while still matching the moments of U.S. consumption growth). Overall, our results show that in the absence of concrete knowledge about investors' information (e.g., whether investors have some news about future dividends or how much noise contaminates the dividend process), it becomes more difficult to conclude that the observed volatility in the data is excessive.

The remainder of the paper is organized as follows. Section 2 describes the model and the information setup. Section 3 examines how investor information influences the volatility of the price–dividend ratio. Section 4 extends the analysis to

[3] We employ a standard unobserved-component time series model for dividend growth. Veronesi (2000) considers a Markov switching process where investors receive a noisy signal about the drift parameter for dividends which can take on different values.

[4] See, for example, Barsky and Sims (2011).

[5] On page 41, West (1988b) acknowledges that his result "may not extend immediately if logarithms or logarithmic differences are required to induce stationarity [of the dividend process]."