The Volatility of Stock Market Prices

ROBERT J. SHILLER

If the volatility of stock market prices is to be understood in terms of the efficient markets hypothesis, then there should be evidence that true investment value changes through time sufficiently to justify the price changes. Three indicators of change in true investment value of the aggregate stock market in the United States from 1871 to 1986 are considered: changes in dividends, in real interest rates, and in a direct measure of intertemporal marginal rates of substitution. Although there are some ambiguities in interpreting the evidence, dividend changes appear to contribute very little toward justifying the observed historical volatility of stock prices. The other indicators contribute some, but still most of the volatility of stock market prices appears unexplained.

WHY ARE STOCK MARKET PRICES SO VOLATILE? THE standard deviation from 1871 to 1986 of the January-to-January percentage change in the real Standard and Poor composite stock price index P_t (Fig. 1, solid line) is 17%. The real index rose 85% between 1927 and 1929, and fell 52% between 1929 and 1932. It rose 69% between 1954 and 1957. It fell 56% between 1973 and 1975. What is it that is so different about the demand for, or supply of, corporate shares from one year to the next that might account for such big price movements?

Price changes have long been attributed to psychological factors: investor overreaction to earnings, dividends, or other news; waves of social optimism or pessimism; fashions or fads. However, statistical evidence amassed during the past few decades has been widely interpreted as implying that markets are "efficient" (*1, 2*).

According to the efficient markets hypothesis, price changes occur when new information about the true investment value of stocks becomes available to the public: the price changes are big because the information is about something very important. However, statistical analyses (*3–9*) of aggregate historical data have recently raised questions of whether something sufficiently large does indeed happen to true investment value to justify the price movements.

The Simple Efficient Markets Model

The simple efficient markets model states that the real return R_t on stocks cannot be forecasted, that all information about future prices is efficiently incorporated in today's price, and one can never know that stocks are a better or worse investment today than at any other time. Formally, the model can be written:

$$E_t R_t = r \tag{1}$$

$$R_t \equiv (P_{t+1} - P_t + D_t)/P_t \tag{2}$$

Here, R_t denotes the total real return (capital gain plus dividend income) from holding stocks between the time t and time $t + 1$ as a proportion of the initial investment P_t; D_t denotes the real dividends paid on the stock between time t and $t + 1$, and r is a constant. E_t denotes the mathematical expectation (or optimal forecast) conditional on the set of information available to the public at time t, a set which includes $P_t, P_{t-1}, P_{t-2}, \ldots, D_{t-1}, D_{t-2}, \ldots$ and current and lagged values of other data. Equation 1 may be described as a random walk hypothesis for prices corrected for dividend payments.

It follows (*10*) from Eq. 1 (and a condition that price does not have an upward trend through time if dividends do not) that price is given by

$$P_t = E_t P_t^* \tag{3}$$

$$P_t^* \equiv D_t/(1 + r) + D_{t+1}/(1 + r)^2 + D_{t+2}/(1 + r)^3 + \ldots \tag{4}$$

That Eq. 3 also implies Eq. 1 can be easily verified by substitution into Eq. 2 (*11*). In words, P_t is the optimal forecast of P_t^*, the true investment value, and P_t^* is the present value, discounted at constant rate r, of actual future real dividends. P_t^* may also be described as the perfect-foresight price, which is the price that would obtain by the efficient markets model (Eqs. 1 or 3) if everyone knew all future dividends with certainty.

The Variability of Forecasts and Forecasted Variables

The optimal forecast of any random variable x cannot be as volatile as x itself unless the forecast is very accurate. If, for example, the forecast of x were as volatile as x but only weakly correlated with it, then high values of the forecast would tend to be associated with negative forecast errors, low values with positive forecast errors. This would then mean that the forecast error was somewhat forecastable, and so the forecast could not be optimal.

Formally, the optimal forecast $E_t x$ must satisfy

$$\sigma(E_t x) = \rho(E_t x, x)\sigma(x) \tag{5}$$

where ρ denotes the correlation coefficient and σ the standard deviation. Since ρ cannot exceed 1.00, in no cases can the volatility of the optimal forecast exceed that of the variable forecasted.

The Variability of Prices and Perfect-Foresight Prices

One can use Eq. 5 to test the model Eq. 3 by computing P_t^* and comparing this with the actual real price P_t. This was done (Fig. 1, dotted line) with the actual real dividend series D_t, which is shown in Fig. 2. The discount rate r was taken as the average real return R_t for this index over the entire sample, which was 8.2% (with a standard error of ±1.6%). Of course no one now knows annual dividends after 1985; it was assumed that these are such as to make

The author is professor of economics, Cowles Foundation, and School of Management, Yale University, New Haven, CT 06520-2125.

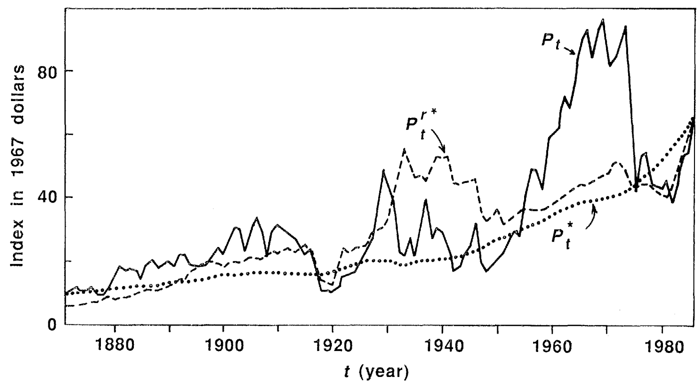

Fig. 1. Real stock prices and their perfect-foresight counterparts, 1871 to 1986 (*4, 9*). P_t (solid line) is the Standard and Poor composite stock price index in January of year t times 100 and divided by π_t, the producer price index for year t. The series π_t is for January starting in 1900, annual average before that. P_t^* (dotted line) is computed from Eq. 4 in the text using the dividend series D_t that is shown in Fig. 2. P_t^{r*} (dashed line) is computed from Eq. 7 in the text using $r_t = \pi_t[\pi_{t+1}\ (1 - r_{1t}/200)\ (1 - r_{2t}/200)]^{-1} - 1$, where r_{1t} is the January value and r_{2t} is the July value of the prime 4- to 6-month commercial paper rate (6-month rate starting in 1979) in annual percent.

$P^*_{1986} = P_{1986}$. Changing this assumption for P^*_{1986} would have only the effect of dragging up or down the recent values of P_t^* [note that $1/(1.082)^{10} = 0.45$, so that the weight the assumption has in determining P^*_{1976} is less than half that in determining P^*_{1986}; that in determining P^*_{1946} is only 4% of that in determining P^*_{1986}].

As can be seen from Fig. 1, the P_t^* so computed is dramatically less variable than the actual real price P_t. The forecast P_t has shown tremendous variability, whereas the variable forecasted, P_t^*, has remained quite close to a smooth trend line. This appears to contradict Eq. 5 where $\sigma(x)$ is the standard deviation of detrended P_t^* (defined as the ratio of P_t^* to an exponential curve, the long-run deterministic "trend," drawn through P^*) and $\sigma(E_t x)$ is the standard deviation of detrended P_t (defined as the ratio of P_t to the same smooth curve). This finding and analogous findings in terms of other measures of variability have been interpreted as seriously calling into question the simple efficient markets model (*3–9*).

One must be cautious in concluding from this evidence that stock prices are too volatile to satisfy Eq. 3. Since the half-life of the weighting pattern defining P_t^* in Eq. 4 is about a decade, then even though we have 116 annual observations, we essentially have no more than about a half dozen more or less independent observations of P_t^*. We should not regard the year-to-year choppiness of P_t when compared to P_t^* as evidence against the model Eq. 1 (*12*). Still, a half dozen observations could suffice to reject a statistical hypothesis convincingly, if the evidence is strong enough. Moreover, although the efficient markets model does not necessarily imply this, it might have turned out that people have a great deal of information about future dividends, which would imply substantial correlation even between the year-to-year changes in P_t and P_t^*; that would be convincing evidence for the efficient markets model, which was not found.

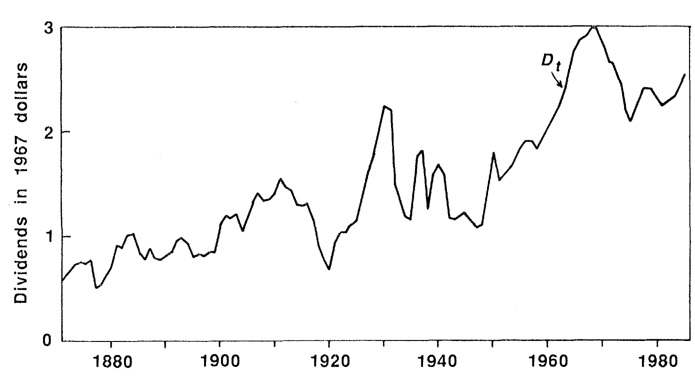

Fig. 2. Real dividends, 1871 to 1985. The nominal dividend series starting in 1926 is dividends per share adjusted to index, four-quarter total, fourth quarter, from Standard and Poor statistical service. The nominal dividends before 1926 are equal to Cowles series Da-1 (*24*) times 0.13 to correct for change in base year in the index. The Cowles series was constructed to correspond to the Standard and Poor composite index (by its earlier base year). D_t is the nominal dividend divided by the annual average producer price index.

Criticisms

The conclusion that the behavior of P_t^* is inconsistent with the efficient markets model Eq. 3 has been quite controversial. One concern of critics (*13*) has related to the fact that with "smooth" time series, the sample standard deviation is in small samples a downward biased measure of the true standard deviation. Since P_t^* is a smoothed version of dividends, there is reason to suspect that the downward bias might be greater for the estimated standard deviation of P_t^* than it is for the estimated standard deviation of P. If P_t were only a little more variable than P_t^*, then we might attribute the apparent violation of Eq. 5 to such bias, and not to the failure of the efficient markets model. However, Kleidon (*14*) argued that this bias alone is unlikely to account for the fact that the standard deviation of detrended P_t is so much higher than that of P_t^*. Another criticism (*15, 16*) of the conclusion against market efficiency is that it relies on the assumption that the trend path of dividends was known in advance: the so-called "deterministic trend" may in fact be a random walk and the uncertainty about future dividends may well have been much larger than is suggested by the slight variations around the ultimate realization of the random walk. This criticism is potentially important but not conclusive (*17*), and several papers have provided alternative tests that sought to deal with this criticism, tests that were interpreted as showing excess volatility relative to Eq. 3 for stock prices (*7–9*).

Price-Dividend Ratios

A simple way to respond to the criticism concerning trend is to detrend by dividing by D_{t-1}, and compare P_t/D_{t-1} with P_t^*/D_{t-1}. Since D_{t-1} is known at time t, the detrending at time t makes use only of information available at that time. Equation 3 implies that P_t/D_{t-1} equals the optimal forecast of P_t^*/D_{t-1}, and P_t^*/D_{t-1} is proportional to a weighted average over k of D_{t+k}/D_{t-1}. The model thus says that P_t/D_{t-1} should be high when dividends can be forecasted to increase in the not-too-distant future, and low when dividends can be forecasted to decrease.

A plot of P_t/D_{t-1} and P_t^*/D_{t-1} appears in Fig. 3. Here, P_t/D_{t-1} is more variable than P_t^*/D_{t-1}, but it is only slightly more variable. However, since the correlation ρ between P_t^*/D_{t-1} and P_t/D_{t-1} is only 0.03, we may say using Eq. 5 with $x = P_t^*/D_{t-1}$ that P_t/D_{t-1} is vastly more variable than it should be given the information it conveys about future dividend changes. It is worth noting, though, that P_t/D_{t-1} and P_t^*/D_{t-1} in Fig. 1 share some movements, each over a few years or so, until the early 1950's. (The correlation between the two series 1872 to 1950 is 0.06.) These shared movements might be regarded as suggesting some element of truth to the efficient markets model (*9*). The shared movements occur because some of the transient short-run movements in dividends are not fully reflected in price, and both P_t/D_{t-1} and P_t^*/D_{t-1} share the same denominator. But this is faint praise for the efficient markets model. If P_t had instead been equal each year just to a simple sum of real dividends over the preceding 20 years, implying P_t had been much less volatile than it was, then the correlation for 1891 to 1986 between P_t/D_{t-1} and P_t^*/D_{t-1} would have been 0.72.

An Alternative Efficient Markets Model with Interest Rate Data

The evidence regarding market efficiency does not take into account any expected variability of real interest rates. To those accustomed to hearing that the market rally since 1982 is due to declining interest rates, this omission will seem important. Equation 1 may be replaced by $E_t(R_t - r_t) = g$, where r_t is the 1-year real interest rate (the nominal interest rate corrected for the actual inflation between time t and time $t+1$), and g is a constant. This alternative acknowledges that stocks earn more on average than do interest-bearing assets, presumably to compensate stockholders for the greater risk inherent in stocks, but that the relative attractiveness of stocks versus interest-bearing assets does not change through time. It follows as above then (disregarding the fact that the 1-year real interest rate r_t is not known with certainty at time t) that

$$P_t = E_t P_t^{r*} \qquad (6)$$

$$P_t^{r*} \equiv D_t/(1 + r_t + g) + D_{t+1}/[(1 + r_t + g)\ (1 + r_{t+1} + g)] + D_{t+2}/[(1 + r_t + g)\ (1 + r_{t+1} + g)\ (1 + r_{t+2} + g)] + \ldots \qquad (7)$$

P_t is the optimal forecast of P_t^{r*}, the true investment value in this model, and P_t^{r*} is just the present value of future dividends discounted by the actual future real interest rates.

As above, P_t^{r*} was computed with data on actual dividends paid and also data on real interest rates, the latter computed with historical data on prime commercial paper. The expected return differential g between stocks and commercial paper was estimated as the sample mean differential from 1871 to 1985: 4.8% with a standard error of $\pm 1.7\%$. The P_t^{r*} with this g and with a terminal condition that makes $P^{r*}_{1986} = P_{1986}$ is plotted in Fig. 1 (dashed line). P_t^{r*} is slightly less variable than P_t, contrary to the implications of the model, and shows virtually no correlation with P_t except for trend. The stock market rally from 1982 to 1986 is indeed explained by the change in real interest rates, but other movements in P_t are not matched by movements in P_t^{r*}. There is a conspicuous shared movement in P_t^{r*} and P_t between the years 1920 and 1929, but a sharply opposite movement between 1929 and 1933. P_t/D_{t-1} and P_t^{r*}/D_{t-1} (Fig. 3) are negatively correlated (1872 to 1986, $\rho = -0.02$) (*18*).

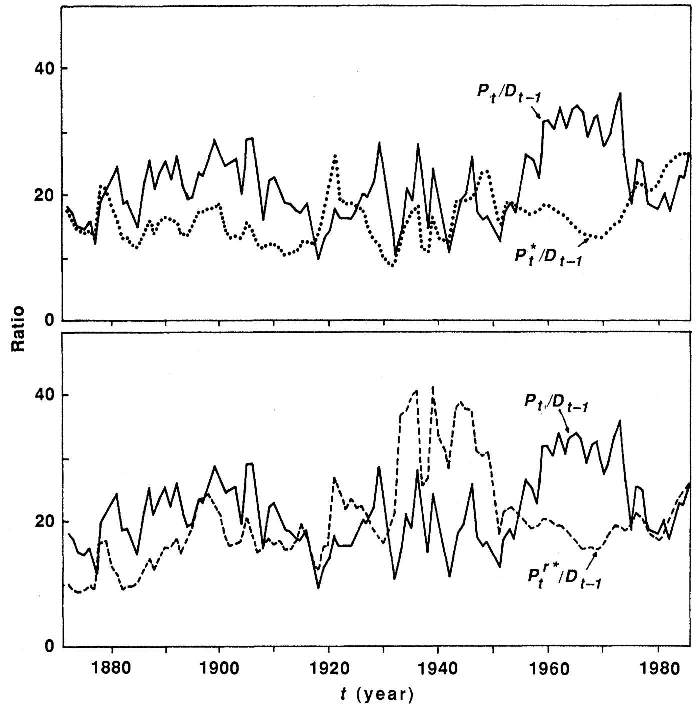

Fig. 3. Price dividend ratio and perfect-foresight counterparts. Shown are series plotted in Fig. 1 divided by the D shown in Fig. 2 for the preceding year.

A Consumption-Based Efficient Markets Model

Another way to incorporate time-varying real discount rates into the efficient markets hypothesis is to look at nonfinancial evidence that might suggest movements in rates of discount. Per capita real-consumption expenditure at time t, C_t, is an indicator of current economic well-being for individual investors. When C_t is low relative to expected future C_t, it is plausible that demand for stocks would tend to be low, and thus expected return (and hence the rate at which future dividends will be discounted to today's price) must be high to induce investors to hold existing shares. In simple terms, in bad times people will be inclined to sell some of the shares to consume the proceeds, and so price must fall to clear the market. For this purpose, bad times must be defined as times when people expect higher C_t in the future; they would not be tempted to consume their savings if they expected worse times yet in the future.

This intuitive notion may be formalized by specifying that the representative investor maximizes an expected utility function $E_tU_t \equiv E_t[u(C_t)/(1+h) + u(C_{t+1})/(1+h)^2 + u(C_{t+2})/(1+h)^3 + \ldots]$ where $u(C_{t+k})$ is the instantaneous utility from consuming C_{t+k}, $k \geq 0$, and h is the subjective rate of time discount. If the representative consumer chooses consumption subject to a budget constraint to maximize this utility function, it follows that the Euler equation $u'(C_t) = E_t[(1 + R_t)u'(C_{t+1})]/(1+h)$ must hold, where u' is the marginal utility, that is, the derivative of the instantaneous utility function. The left-hand side of the Euler equation is the rate at which marginal utility is given up by consuming an increment less today. The right-hand side is the rate at which expected marginal utility is gained by consuming the proceeds of the investment the following period. If the former did not equal the latter, the individual would be better off either by saving more or consuming more. It follows from the Euler equation as above that P_t is given by

$$P_t = E_t P_t^{c*} \qquad (8)$$

$$P_t^{c*} \equiv D_t S_{t1} + D_{t+1} S_{t2} + D_{t+2} S_{t3} + \ldots \qquad (9)$$

$$S_{tk} \equiv u'(C_{t+k})/[u'(C_t)\ (1+h)^k] \qquad (10)$$

Here, P_t^{c*} is the true investment value and S_{tk} is the intertemporal marginal rate of substitution between C_t and C_{t+k}.

We can compute P_t^{c*} using historical data on C_t and D_t if we assume a functional form for $u(C_t)$; a common form in the theoretical finance literature is $u(C_t) = C_t^{1-A}/(1-A)$, $A > 0$. The larger the parameter A (the coefficient of relative risk aversion), the more the representative investor is assumed to dislike variation in consumption, and the more volatile the P_t^{c*} will be. With $A = 4$, and $1 + h$ chosen as the sample mean for 1889 through 1984 of $(1 + R_t)\ (C_t/C_{t+1})^4$, P_t^{c*} is about as volatile as P_t (Fig. 4) (*6*).

For the earlier part of the sample there is some positive correlation in P_t^{c*} and P_t beyond trend, although there is essentially no correlation in the years since 1950. The correlation we have discovered here is essentially the well-known correlation of the stock

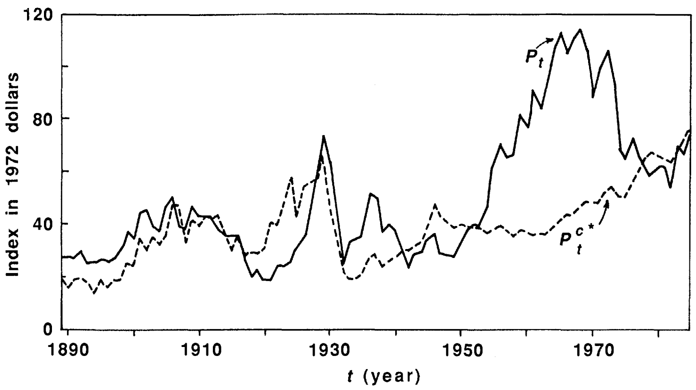

Fig. 4. Real stock prices and perfect-foresight counterpart, 1889 to 1985 (*6, 20*). P_t (solid line) is the annual average Standard and Poor composite stock price index times 100 divided by π'_t, the consumption deflator for nondurables and services for the year. P^{c*}_t (dashed line) is computed from Eq. 9 in the text using the nominal dividend series described in the legend to Fig. 2 (*24*) times 100 and divided by π'_t. C_t is the total for the year real per capita consumption on nondurables and services.

market with the so-called "business cycle," a correlation that has diminished since 1950 (*6*).

For large enough A, (considerably above 4) we will find that Eq. 5 is satisfied where x is detrended P^{c*}_t. Grossman, Melino, and Shiller (*19*) estimated $A = 21.2$ for the sample 1890 to 1983. Such a high value of A may be implausible. For this estimated A, the Euler equation implies, for example, that a person who knows that consumption C_t will increase 5% between this year and next will require an expected return on stocks of over 180% in order to be willing to hold the existing shares.

Anyway, movements in the aggregate stock market are largely independent of movements in prices of other long-term assets such as housing or land (*20*). This is not what we would expect if these rate of discount movements dominated the movements in the stock market.

Statistical Issues

Changes in stock market prices have a probability density that departs distinctly from the usual bell-shaped normal density. The density is "fat-tailed," meaning that occasional big outlier observations tend to occur, outliers that are many sample standard deviations from the mean. The stock market tends to bring up occasional new "surprises" that seem outside the expected range provided by conventional statistical methods.

The fact that the usual normal density is not really well suited for statistical analysis poses serious problems for formal hypothesis testing with financial data. Mandelbrot (*21*) has suggested that the normality assumption be replaced by the assumption of another distribution in the stable Paretian family. Such distributions have infinite standard deviations. The sample standard deviation will of course always be finite, but in no way can it be regarded as an estimate of the true standard deviation. In this case, one cannot use Eq. 5, which was the foundation of the above analysis.

Our understanding of the fat-tailed nature of changes in stock prices has been altered by the discovery that the variability of changes in stock prices changes through time and that this variability can be forecasted. For example, Bollerslev (*22*) has shown in the context of a "generalized auto-regressive conditional heteroskedasticity model" of time-varying variability that the distribution of changes in the monthly postwar Standard and Poor composite index conditional on recent changes is somewhat less fat-tailed than the unconditional distribution. In his model, the unconditional probability density of stock price changes does have a finite standard deviation.

The problem posed by the fat-tailed density is however just one manifestation of a deeper statistical problem. There is no well-defined theoretical statistical model of true investment value, and thus there is no way to be sure that any measures of historical variability accurately represent the potential variability. It is conceivable, for example, that the stock market's ups and downs during the last century represented genuine new information about a potential big disaster (for example, a nationalization) or big windfall (for example, a technological revolution) that would utterly change the outlook for future dividends. There is really no way to ensure that people were not right to change their minds from time to time about the possibility of such a rare big event. All we can say is that nothing actually happened in history that would seem to justify the price movements. Those who criticize the studies finding excess volatility in stock prices under the assumption that dividends are a random walk (*14*) or that managers smooth dividends (*15*) are really also relying on an assumption that the outlook for future dividends is much different than the historical record would suggest: their criticisms imply that dividends will not always stay as close to a simple trend line in the future as they have in the past.

Conclusion

The price P_t or its ratio to dividend P_t/D_{t-1} generally appears to show too much variability given its correlation with its perfect-foresight counterpart under any of the models considered here. This is not to say that there might not be an element of truth to some of the models (witness some of the shared movements in subperiods of the series noted above), and statistical significance of these results is a difficult issue still unresolved in the literature. But there is certainly little indication that the source of stock price movements ought to be considered explained by any of these models.

For the aggregate stock market, the widespread impression that there is strong evidence for market efficiency may be due just to a lack of appreciation of the low power of many statistical tests (*23*). It should be borne in mind, however, that there are individual investment assets whose true investment value does not look like a simple trend; for some of these, true investment value may predictably change sharply, even by orders of magnitude. For such assets, the efficient markets hypothesis does appear to suggest useful models. The notion of efficient markets, of course, also has value in the simple sense that stock market returns are not highly forecastable.

REFERENCES AND NOTES

1. E. Fama, *J. Financ.* **25**, 383 (1970).
2. B. Malkiel, *A Random Walk Down Wall Street* (Norton, New York, 1985).
3. S. LeRoy and R. Porter, *Econometrica* **49**, 555 (1981).
4. R. Shiller, *Am. Econ. Rev.* 71, 421 (1981).
5. ———, *J. Financ.* **36**, 291 (1981).
6. S. Grossman and R. Shiller, *Am. Econ. Rev.* **71**, 222 (1981).
7. N. G. Mankiw, D. Romer, M. Shapiro, *J. Financ.* **40**, 677 (1985).
8. K. West, "Speculative bubbles and stock price volatility" (Working paper 1833, National Bureau of Economic Research, Cambridge, MA, 1986).
9. J. Campbell and R. Shiller, "Cointegration and tests of present value models" (Working paper 1885, National Bureau of Economic Research, Cambridge, MA, 1986).
10. P. Samuelson, *Ind. Manage. Rev.* **6**, 41 (1965).
11. Two properties of the conditional expectations operator E_t are used here: $E_sE_t = E_{\min(s,t)}$ and $xE_t = E_tx$, where x is any random variable in the conditioning information set at time t.
12. R. Shiller, *J. Polit. Econ.* 87, 1190 (1979).
13. See, for example, M. Flavin, *ibid.* **91**, 929 (1983).
14. A. Kleidon, *J. Bus.* **59**, 237 (1986).
15. ———, *J. Polit. Econ.* **94**, 953 (1986).
16. T. Marsh and R. Merton, *Am. Econ. Rev.* **76**, 483 (1986).

17. R. Shiller, *ibid.*, p. 499.
18. J. Campbell and R. Shiller, unpublished manuscript.
19. S. Grossman, A. Melino, R. Shiller, "Estimating the continuous time consumption based asset pricing model" (Working paper 1643, National Bureau of Economic Research, Cambridge, MA, 1985).
20. R. Shiller, *Carnegie Rochester Conf. Ser. Public Policy* 17, 203 (1982).
21. B. Mandelbrot, *J. Bus.* 36, 394 (1963).
22. T. Bollerslev, "A conditionally heteroskedastic time series model for security prices and rates of return data" (discussion paper 85-32, Department of Economics, University of California, San Diego, 1985).
23. L. Summers, *J. Financ.* 41, 591 (1986).
24. Alfred Cowles & Associates, *Common Stock Indexes* (Principia, Bloomington, IN, ed. 2, 1939).
25. Supported by NSF grant SES-8520075.

Laser Remote Sensing of the Atmosphere

DENNIS K. KILLINGER AND NORMAN MENYUK

Laser beams can be used as long-range spectroscopic probes of the chemical composition and physical state of the atmosphere. The spectroscopic, optical, and laser requirements for atmospheric laser remote sensing are reviewed, and the sensitivity and limitations of the technique are described. A sampling of recent measurements includes the detection of urban air pollution and toxic chemicals in the atmosphere, the measurement of global circulation of volcanic ash in the upper atmosphere, and the observation of wind shear near airports.

LASER REMOTE SENSING OF ATMOSPHERIC PROPERTIES from a single location is referred to as lidar, an acronym for light detection and ranging, and is analogous to radar (*1–6*). In lidar, the projection of a short laser pulse is followed by reception of a portion of the radiation reflected from a distant target or from atmospheric constituents such as molecules, aerosols, clouds, or dust (Fig. 1). The incident laser radiation interacts with these constituents, causing alterations in the intensity and wavelength according to the strength of this optical interaction and the concentration of the interacting species in the atmosphere. Consequently, information on the composition and physical state of the atmosphere can be deduced from the lidar data. In addition, the range to the interacting species can be determined from the temporal delay of the backscattered radiation. Lidar has been used to measure the movement and concentration of air pollution near urban centers, the chemical emission around industrial plants, and atmospheric trace chemicals in the stratosphere. Lidar has also been used to measure the velocity and direction of winds near storms and airports and to track the global circulation of volcanic ash emitted into the atmosphere after recent eruptions, such as those at Mount St. Helens and El Chichón. Under optimal conditions lidar can be extremely sensitive. An example is the ground-based laser remote sensing of sodium and lithium atoms in the stratosphere at ranges greater than 90 km and in concentrations as low as a few atoms per cubic centimeter (*5*). More commonly, detection ranges are on the order of a few hundred meters to several kilometers and concentration levels on the order of parts per million to parts per billion.

The use of optical backscatter to measure properties of the atmosphere is not new, extensive experiments having been conducted in the early 1900's with large searchlights. The field of optical remote sensing was greatly advanced by the laser, which offers several improvements over conventional light sources. These include narrow spectral width (<0.01 nm), a frequency or color that is often tunable, and high peak power ($>10^6$ W) available in a short pulse (<1 μsec) and in a narrow beam (<10 cm in diameter). These attributes make the laser an ideal spectroscopic probe of the atmosphere. In this regard, a lidar system may be thought of as an "active" remote sensing system since it can illuminate the target region, in contrast to a "passive" optical sensor which detects ambient light or thermal emission from the target.

The applications of lidar systems for the remote sensing of atmospheric properties were appreciated soon after the discovery of lasers in the early 1960's. Early lidar measurements were made in 1962 by Fiocco and Smullin (*7*) who bounced a laser beam off the moon and who also investigated the turbid layers in the upper atmosphere; in 1963 by Ligda (*8*) who used a ruby laser to obtain the first lidar measurements of cloud heights and tropospheric aerosols; and in 1964 by Schotland (*9*) who used a temperature-tuned ruby laser to detect water vapor in the atmosphere. Progress has been continuous since that time, but the discovery of different laser sources in the past decade, coupled to improvements in optical instrumentation and data processing, has been responsible for the recent surge in the number of laser remote sensing systems. This improved capability has been accompanied by an increased awareness of the need to monitor the impact of natural and anthropogenic influences on the environment.

Remote sensing of the atmosphere by optical techniques can be accomplished in several ways. One technique involves measurement of the absorption spectrum of the atmosphere over a long path separating a spectroscopic optical source and a detector (*5*). Another long-path absorption technique uses a configuration, described by Hinkley and Kelley (*10*), in which a tunable laser source and detector are located together and a retroreflecting mirror is placed at a distance of several hundred meters; such a system is useful when the laser source is weak, since the retroreflector greatly enhances the returned radiation. In this article we will primarily be concerned with a pulsed lidar system where the laser and detector are located together and no retroreflector is used as the target; in this case the returned laser radiation is due to backscatter from aerosols or dust in the atmosphere or a topographic target such as a hill or trees.

The authors are at the Massachusetts Institute of Technology, Lincoln Laboratory, Lexington, MA 02173-0073.