

A New Dividend Forecasting Procedure That Rejects Bubbles in Asset Prices: The Case of 1929's Stock Crash

R. Glen Donaldson

University of British Columbia

Mark Kamstra

Simon Fraser University

We develop a new procedure to forecast future cash flows from a financial asset and then use the present value of our cash flow forecasts to calculate the asset's fundamental price. As an example, we construct a nonlinear ARMA-ARCH-Artificial Neural Network model to obtain out-of-sample dividend forecasts for 1920 and beyond, using only in-sample dividend data. The present value of our forecasted dividends yield fundamental prices that reproduce the magnitude, timing, and time-series behavior of the boom and crash in 1929 stock prices. We therefore reject the popular claim that the 1920s stock market contained a bubble.

Many empirical tests of asset price behavior call for the comparison of an asset's market price to its

This article is a substantial revision of our earlier paper, "Using Dividend Forecasting Models to Reject Bubbles in Asset Prices: The Case of 1929's Stock Crash," which received the New York Stock Exchange prize for best paper on equity trading presented at the 1994 Western Finance Association meetings. In revising this article we received numerous useful suggestions from a great many of our colleagues. Extended discussions and correspondence with Franklin Allen (the editor), John Campbell, Burton Hollifield, Allan Kleidon, Lisa Kramer, Alan Kraus, Rick Mishkin, Gregor Smith, and an anonymous referee of this journal were particularly helpful. We also gratefully acknowledge financial support from the Social Sciences and Humanities Research Council of Canada. We are responsible for any remaining errors, omissions, or opinions. Please address correspondence to R. Glen Donaldson, Faculty of Commerce and Business Administration, University of British Columbia, 2053 Main Mall, Vancouver, BC V6T 1Z2, Canada.

The Review of Financial Studies Summer 1996 Vol. 9, No. 2, pp. 333-383
© 1996 The Review of Financial Studies 0893-9454/96/\$1.50

"fundamental" price, defined as the market's expected discounted present value of future cash flows. Since we cannot observe investors' true expectations, however, the market price is in practice compared to the econometrician's *estimate* of the fundamental price. Traditional fundamentals estimation procedures often involve either replacing expected cash flows with actual cash flows, as in Shiller's (1981) seminal study, or assuming that market participants expect future cash flows to grow at some constant rate, as in the popular Gordon (1962) model.

There is a vast and growing literature that suggests that the prices produced by traditional fundamentals estimation procedures do not share important properties with actual stock price data. Perhaps the most dramatic example is the widely perceived failure of traditional models to provide fundamental explanations for the type of extreme price fluctuations observed during supposed "bubble episodes," such as the Great Stock Market Crash of 1929, the South Sea Bubble and the Mississippi Bubble.¹ Indeed, the difficulty in explaining these rare but important events has led some economists to assert that asset price movements are influenced by factors other than cash flow fundamentals, the assumption being that the traditional fundamentals estimation procedures being employed are sound but that the standard present value model is at odds with the data. However, there are other economists who take the opposite approach and argue instead that the standard present value model is sound but that misspecifications in traditional fundamentals estimation procedures give the false impression of bubbles in asset prices where there are in fact none.² Thus, while the first point of view implies that some types of price behavior can never be explained by fundamentals, the second viewpoint argues that, if we could only find a well-specified fundamentals estimation procedure, then the standard present value model could explain even the most extreme price fluctuations as fundamental events.

Given the preceding discussion, the purpose of our article is to develop a new methodology for calculating fundamental asset prices. In particular, we develop a new—and we argue more accurate—procedure to forecast future cash flows from a financial asset and to then use the expected present value of these forecasted future cash flows to estimate the asset's fundamental price. Since traditional fundamentals estimation procedures seem particularly challenged by

¹ For a discussion of these famous historical events, and a summary of the evidence both for and against the claim that they reveal bubbles in asset prices, see Garber (1990), Kindleberger (1978), Shiller (1989), and White (1990).

² See, for example, Flood and Garber (1994), Flood and Hodrick (1990), and Kleidon (1986). Akerl and Smith (1993) also raise questions about the completeness of cash flow data traditionally employed in volatility studies.

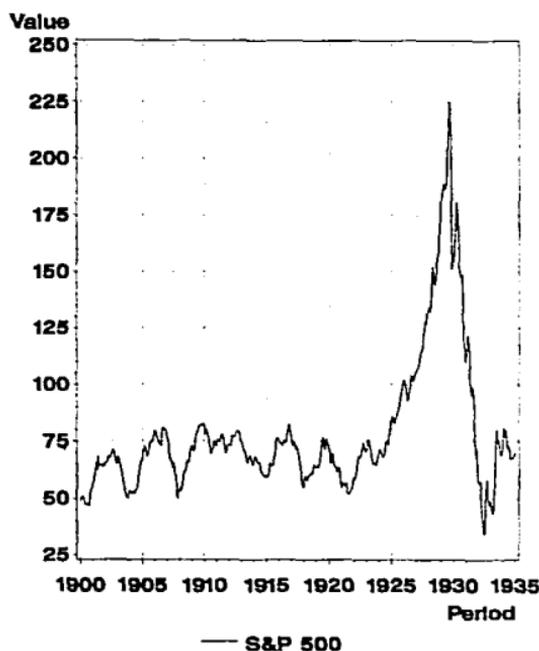


Figure 1

S&P 500 stock price index

This figure plots the S&P 500 stock price index. See Appendix 1 for data sources.

extreme events, and since the stock boom and crash of 1929 is the most widely cited example of a stock price bubble for which we have reliable data, we use the panic of 1929 as a laboratory in which to compare our new procedure to traditionally employed models.

The extreme nature of the 1920s stock boom and crash is evident in Figure 1, which plots the S&P 500 stock price index monthly from 1900 to 1934 (data sources are listed in Appendix 1). Notice that, from 1900 to 1925, the S&P 500 mostly assumed values within the fairly narrow range of 50 to 80 points and returned to this range after 1932. Between 1926 and 1932, however, the index rose rapidly to peak at 225 in September 1929 and then crashed with the infamous October panic.

White's (1990) summation of existing literature on the crash of 1929 leads him to conclude that the "conventional wisdom" and "most commonly accepted version of the boom and crash" is that an explosive "bubble" pushed the market price well above the fundamental price during the late 1920s and that this bubble burst with the 1929 panic. Our new fundamentals estimation procedure yields the opposite result. In particular, our econometric procedure—in conjunction with the present value model—yields fundamental prices that reproduce

the magnitude, timing, and time-series behavior of the boom and crash in 1929 stock prices. We therefore reject the popular claim that the 1920s stock market contained a bubble and instead argue that, at least in this one case, the standard bubble-free present value model remains a valid tool for explaining market behavior.

Our procedure for obtaining fundamental prices is as follows. First, we use dividend data from before 1920 to estimate a nonlinear ARMA-ARCH model for the time-series behavior of discounted dividend growth. Second, we use our model, with in-sample data, in a Monte Carlo experiment to produce out-of-sample forecasts of discounted dividend growth for 10,000 different simulated economies into the (almost) infinite future. Third, we calculate the present discounted value of each of our 10,000 different forecasted dividend streams to obtain 10,000 different possible prices. Finally, the cross-sectional mean of these 10,000 simulated prices is computed to obtain our estimate of the market's expected discounted present value of the asset's future cash flows, and thus the fundamental price.

The purely out-of-sample dividend forecasts from our model are able to produce a fundamental price series that rises from roughly the same value as the market price in 1920 to peak at the same time, and within 10 points of, the market price in 1929 and then crashes along with the market price through the early 1930s. Furthermore, on comparing our fundamental prices to market prices with a number of statistical tests, we are easily able to reject the hypothesis that market prices contain a bubble. Indeed, our dividend forecasts suggest that, given the information available to market participants living in the early 1920s, dividends were expected to increase by enough to warrant the observed rise in stock prices. However, as new information arrived in the late 1920s, expectations of future dividends were revised downward, resulting in the observed crash in prices. We therefore conclude that there was *not* a bubble in 1920s stock prices and thus that the standard bubble-free present value model is appropriate for even this most extreme episode.

The remainder of our article proceeds as follows. In Section 1 we formalize the terms "fundamental price" and "bubble" and demonstrate that traditional procedures for estimating fundamental prices erroneously find bubbles in asset prices because traditional models are misspecified. In Section 2 we present our nonlinear ARMA-ARCH model and Monte Carlo simulation procedure for forecasting dividends and demonstrate that, when market prices are compared to the fundamental prices produced by our new procedure, we reject the hypothesis of a bubble in market prices. In Section 3 we provide some insight into why various elements of our econometric representation are important for achieving a well-specified model to accurately fore-

cast dividends and thus why we reject the bubbles hypothesis while traditional models do not. Section 4 concludes.

1. The Data and its Traditional Interpretations

We begin our investigation by comparing market prices to the fundamental prices obtained from traditional models. Such a comparison is useful because it provides a benchmark against which to compare the results of our new fundamentals estimation procedure and because an understanding of why traditional models fail to reject bubbles helps to motivate the particular approach we employ.

1.1 Defining fundamentals and bubbles

Consider a share of stock whose price is determined at the beginning of each period and that pays a dividend at the end of each period. Define D_t as the dividend payment made at the end of period t , P_{t+1} as the stock's selling price at the beginning of period $t + 1$, and $1 + r_t$ as the gross real rate investors use to discount payments received during period t . The rational time t price of the stock is then given by Equation (1),

$$P_t = E_{t-1} \left\{ \frac{D_t + P_{t+1}}{1 + r_t} \right\} \quad (1)$$

in which E_{t-1} is the expectations operator conditional on information available to the market when P_t is being determined at the beginning of period t (i.e., information from the end of period $t - 1$ and earlier). Provided that the discounted present value of the stock's price infinitely far into the future is zero (i.e., there are no bubbles), we can recursively substitute for future prices in Equation (1) to find that the fundamental price, P_t^F , equals the expected present discounted value of all future dividends, as in Equation (2),

$$P_t^F = E_{t-1} \left\{ \sum_{j=0}^{\infty} \left(\prod_{i=0}^j \frac{1}{1 + r_{t+i}} \right) D_{t+j} \right\}. \quad (2)$$

Following authors such as Camerer (1989) and West (1988), a rational explosive bubble is defined as a price process that satisfies Equation (1) but not Equation (2). For example, if P_t^M is the market price and P_t^F is the fundamental from Equation (2), then the market price will satisfy Equation (1) for any time series B such that $P_t^M = P_t^F + B_t$ and $B_t = E_{t-1}\{B_{t+1}/(1 + r_t)\}$; i.e., although investors rationally know that the current market price, P_t^M , exceeds the present value of future dividend payments, P_t^F , the value of the bubble term B_t is expected to increase just fast enough so that the higher than fundamental price

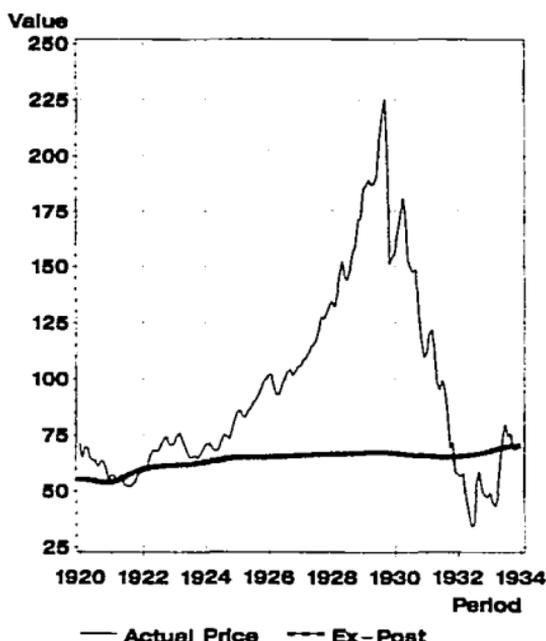


Figure 2
Actual price versus ex-post warranted price
 This figure plots the S&P 500 stock price index (plain line) versus the ex-post warranted price (line joining stars) obtained by substituting into Equation (2) the realized S&P 500 dividends for D_t and the riskfree interest rate plus a constant equity premium for r_t (the actual price as of December 1933 is used for the truncation approximation). All calculations are performed using real variables and then plotted in nominal terms to retain the visual image of the boom and crash. See Appendix 1 for data sources.

an investor pays for the asset today is rewarded by an even higher than fundamental expected selling price next period. Conversely, an irrational bubble is defined as a divergence from fundamental prices that satisfies neither Equation (1) nor Equation (2).

1.2 Ex-post warranted prices

Much of the traditional quantitative evidence in support of bubbles compares the actual market price to the ex-post warranted price, calculated as the discounted present value of future *realized* dividends. Following the procedure of Shiller (1981) and others, the “fundamental” to which the actual price is compared is thus obtained by substituting realized future dividends and discount rates into Equation (2) for D_{t+j} and r_{t+i} , respectively.

Figure 2 plots the actual value of the S&P 500, monthly from 1920 to 1933, along with the ex-post warranted price obtained as the present discounted value of future realized real dividends (data sources in

Appendix 1).³ Visual inspection of Figure 2 suggests that, if investors had perfect foreknowledge of future dividend behavior, then the only way they would have supported the spectacular rise and fall in observed prices as a market equilibrium is if the market price contained an explosive bubble.

Our finding from visual examination of Figure 2 is supported by Table 1, which presents some summary statistics on prices and returns as well as unit root tests for the presence of bubbles.⁴ Column 1 lists the information reported in each row. Results for the actual S&P 500 are reported in column 2. Column 3 reports results from the ex post warranted series plotted in Figure 2. Other columns report results for alternative models to be discussed below.

Part A of Table 1 contains some summary statistics on the actual and various model-generated prices. First, note from the first three rows of statistics in Table 1 that the ex post warranted price is roughly 150 points below the actual price on the peak date (September 1929). Second, note that if there is an explosive bubble in prices then the difference between the actual price and the fundamental price (i.e., the price error) will be nonstationary, since the divergence of actual prices from fundamental prices will be persistent in nature.⁵ The last row of part A in Table 1 reports the *t*-statistic from a Phillips-Perron test of the null hypothesis that the real price error follows a unit root process.⁶ For the ex post warranted price we fail to reject the unit root null for the period 1920 to 1933. The same (unreported) result holds over the 1920 to 1929 subperiod. Thus, under the assumption that the ex post model accurately represents market fundamentals (i.e., that the model is well specified), we find evidence for a bubble in the market price.

³ All "fundamentals" calculations are performed using real variables and then translated into nominal values to retain the visual image of the boom and crash. Since we do not have data on realized dividends into the infinite future, we follow the convention adopted by Shiller (1981) and replace the present value of dividends outside our data set with the present value of the asset price on the last date for which we do have data. In Figure 2, the discount rate employed is the real yield on high-quality short-term debt plus a constant equity premium equal to the average return on stocks over short-term low-risk bonds from 1871 to 1988 (see Appendix 1). As has been amply demonstrated in the literature, discounting dividends at a constant rate, or even at a variable rate determined by the marginal utility of consumption, leads to the same conclusion.

⁴ There are some additional tests that have been used in the bubbles literature but which we do not perform in our article. Our choice of bubbles tests has been dictated by our desire to maintain a level playing field among the many models we investigate. For example, we do not use West's (1987) popular test because both Barsky and DeLong's (1993) augmentation of the Gordon model and our own nonlinearly augmented ARMA-ARCH forecasting model do not deliver the overidentifying restrictions necessary to implement the test.

⁵ See Campbell and Shiller (1987) for a discussion of the time-series relationship between market prices and dividend-based fundamentals.

⁶ Results reported are from unit root tests with data-dependent lag length and no trend. Adding a trend does not affect our conclusions.

Table 1
Summary statistics for fundamental prices and returns, 1920–1933

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Model name	Actual S&P 500 series	Expost warranted series	Basic Gordon model	Augmented Gordon model	Forecast simulation constant- <i>r</i>	Forecast simulation consumption- <i>r</i>	Forecast simulation bond yield- <i>r</i>
Data plotted in:	Figures 1–7	Figure 2	Figure 3	Figure 4	Figure 5	Figure 6	Figure 7
Panel A: Prices							
Peak date	09/29	07/29	03/30	05/30	12/29	05/30	09/29
Peak value	225.19	67.44	114.82	184.92	233.26	244.52	215.72
Value at 09/29	225.19	67.41	111.91	172.47	209.37	160.80	215.72
P/D correlation	—	-0.083	-0.064	0.383*	0.543*	0.318*	0.806*
Unit root <i>t</i> -statistic (real price error)	—	-1.31	-1.99	-1.90	-3.02*	-2.35	-5.21*
Panel B: Returns							
Mean	.00715	.00623	.00220	.00313	.00732	.00675	.00690
Standard deviation	.07237	.01334	.01715	.02690	.08709	.08664	.09874
Third moment	.00055	-4.9E-6	-2.9E-6	-0.00001	.00051	.00029	.00029
Fourth moment	.00045	2.6E-7	3.4E-7	2.0E-6	.00023	.00021	.00036
Unit root <i>t</i> -statistic (real returns)	-9.56*	-1.69	-1.64	-1.58	-12.11*	-11.77*	-13.89*

This table presents summary statistics for the actual S&P 500 and the fundamentals from the various models listed at the top of each column. The first row of Panel A reports the date on which the series listed in each column peaks. The second row gives the value of the series in question on the date of its peak. The third row gives the value of the series in question on September 1929, the month in which actual market prices peak. The fourth row of Panel A reports the correlation between the actual market price-dividend (P/D) ratio and the ratio of the fundamental price in question to dividends (i.e., the fundamental P/D ratio). The last row of Panel A reports the *t*-statistic from a Phillips-Perron test—with data-dependent lag length and no trend—of the null hypothesis that the real price error (i.e., the actual real market price minus the real fundamental price) follows a unit root process. The first four rows of Panel B report the first four moments of the returns distribution from the price series listed at the top of each column. The last row of the table gives the *t*-statistic from a Phillips-Perron test—with data-dependent lag length and no trend—of the null hypothesis that the real return implied by the price series in question possesses a unit root. Throughout the table, * denotes statistical significance at 10 percent.

Part B of Table 1 presents some statistics on the real returns implied by the various price series, including the first four moments of each returns distribution and a test for a unit root in the returns process. As noted by Camerer (1989), if there is a bubble in actual prices then actual returns will have a higher variance and fatter tails than fundamental returns since, as bubbles expand and then burst in the actual series, unusually sizable returns will be realized. Results from the first four rows of part B of Table 1 therefore support the bubbles hypothesis in that the ex post warranted returns series has a much smaller standard deviation and fourth moment than the actual series. The last row of Table 1 presents t -statistics from a Phillips-Perron test of the null hypothesis that returns from the price series in question possess a unit root. The result that actual returns reject the null, while the ex post series does not, is consistent with the assertion of Shiller (1989) and others that returns from the ex post series are much too smooth relative to actual returns.

From Figure 2 and Table 1 it is clear that if ex post warranted prices provided an accurate estimate of the market's fundamentals, then we would fail to reject the bubbles hypothesis. However, we argue that ex post warranted prices do *not* provide the most accurate representation possible of the market's true fundamentals. In particular, a standard assumption that permits us to substitute realized dividends into Equation (2) for their expected values is that dividends are stationary. However, there is substantial evidence that dividends are not stationary and thus that at least one of the assumptions required to treat the ex post warranted series as an accurate representation of the market's fundamental price is violated.⁷ We therefore argue that the estimated fundamentals do not match actual prices in Figure 2 because the estimated fundamentals do not match the market's true fundamentals and not because the true fundamentals do not match actual prices.

1.3 The Gordon growth model

The Gordon (1962) model seeks to address the nonstationarity of dividends by using dividend growth rates, instead of levels. This is accomplished by first rewriting the right-hand side of Equation (2) in terms of the most recently paid dividend D_{t-1} and expected future discounted dividend growth. Define $g_t \equiv (D_t - D_{t-1})/D_{t-1}$ as the growth rate of real dividends, so $D_t = (1 + g_t)D_{t-1}$, to rewrite Equation (2) as $P_t^F = E_{t-1} \left\{ \sum_{j=0}^{\infty} D_{t-1} \Pi_{i=0}^j [(1 + g_{t+i})/(1 + r_{t+i})] \right\}$ which, defining

⁷ Augmented Dicky-Fuller and Phillips-Perron tests for a unit root in our real dividend series cannot reject the unit root null at the 10 percent level of significance. For more on the nonstationarity of dividends, see Kleidon (1986) and Mankiw, Romer, and Shapiro (1985, 1991).

$y_t \equiv (1 + g_t)/(1 + r_t)$ as the discounted real dividend growth rate, can be expressed as $P_t^F = D_{t-1} E_{t-1} \{y_t + y_t y_{t+1} + y_t y_{t+1} y_{t+2} + \dots\}$. Equation (2) can then be rewritten as Equation (3).

$$P_t^F = D_{t-1} E_{t-1} \left\{ \sum_{j=0}^{\infty} \left(\prod_{i=0}^j y_{t+i} \right) \right\}. \quad (3)$$

Gordon (1962) assumes that r (the discount rate) and g (the dividend growth rate) are constants so that $y_{t+i} = y \equiv (1 + g)/(1 + r)$ is a constant for all i in Equation (3). Equation (3) therefore reduces to the familiar Equation (4) in which the fundamental price is a constant multiple of the most recently paid dividend.

$$P_t^F = D_{t-1} \left(\frac{1 + g}{r - g} \right). \quad (4)$$

Using average *monthly* values of r and g from 1900 to 1919, as might an investor living in 1920, produces $(1 + g)/(r - g) = 200$. Figure 3 therefore plots dividends multiplied by 200 (i.e., the basic Gordon model's fundamental price) along with the market price.

White (1990:72) states that, by examining the type of relationship between prices and dividends plotted in Figure 3, we can observe "the remarkable change that overtook the stock market [during the late 1920s]. From 1922 to 1927 dividends and prices moved together, but while dividends continued to grow rather smoothly in 1928 and 1929, stock prices soared far above them," an implication being that there was a bubble in late 1920s stock prices. Indeed, the statistics presented in column 4 of Table 1 for the basic Gordon model give the appearance of a bubble in the market price. For example, as seen in the last row in part A of Table 1 (column 4), we fail to reject the null of a unit root in the difference between the market price and Gordon price, a finding which is consistent with the presence of a nonstationary bubble in the market price. Of course, the conclusion that the market price contains a bubble rests on the assumption that the Gordon model accurately represents the dividend forecasting procedure used by 1920s investors. However, since dividend growth is far from constant in reality—and, more importantly, possesses time-varying conditional moments—we would argue that a model that simply assumes unchanging growth rates is misspecified and thus that the bubble conclusion is unwarranted.⁸

⁸ There is an interesting special case in which Gordon's assumptions would be warranted. If utility is the log of consumption, and if consumption equals dividends, then investors will set $P_t^F = D_{t-1}(\beta/(1 - \beta))$, where β is the standard time discount factor. If the growth rate in dividends is expected to equal Gordon's constant g and the expected return is a constant equal to Gordon's

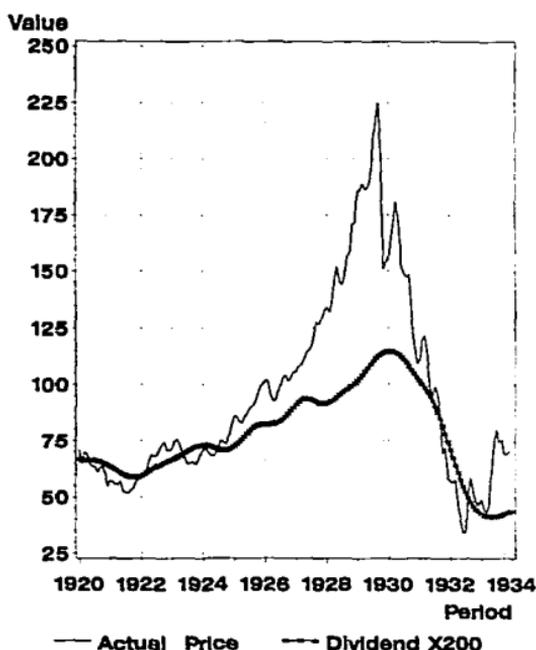


Figure 3
Actual price versus basic Gordon growth fundamental

This figure plots the S&P 500 stock price index (plain line) versus the level of S&P 500 dividends multiplied by 200 (line joining stars). The line joining stars is also the fundamental price implied by the Gordon growth model in Equation (4), since using average *monthly* values of r and g from 1900 to 1919, as might an investor living in 1920, produces $(1 + g)/(r - g) = 200$. See Appendix 1 for data sources.

Recognizing the aforementioned shortcoming in Gordon's original constant- g model, Barsky and DeLong (1993) assume instead that g in Equation (4) is nonconstant and evolves with a geometrically declining distributed lag as in Equation (5), with the weighting parameter λ slightly less than unity,

$$g_t = (1 - \lambda) \sum_{i=0}^{t-1} \lambda^i g_{t-1-i} + \lambda^t g_0. \quad (5)$$

Barsky and DeLong find that they are able to insert Equation (5) into Equation (4) to obtain "fundamental" prices that roughly approximate the broad swings observed in annual stock price data from 1880 to 1992, provided they also assume that $\lambda = 0.97$ and $r = 0.06$.⁹

r , then it is possible to show that $(1 + g)/(r - g) = (\beta/(1 - \beta))$. We thank the referee for drawing this to our attention.

⁹ Barsky and DeLong (1993) also investigate $\lambda = 0.95$, but find that with $\lambda = 0.97$ they are better able to capture broad features of the data. We therefore use $\lambda = 0.97$ in our examples below.

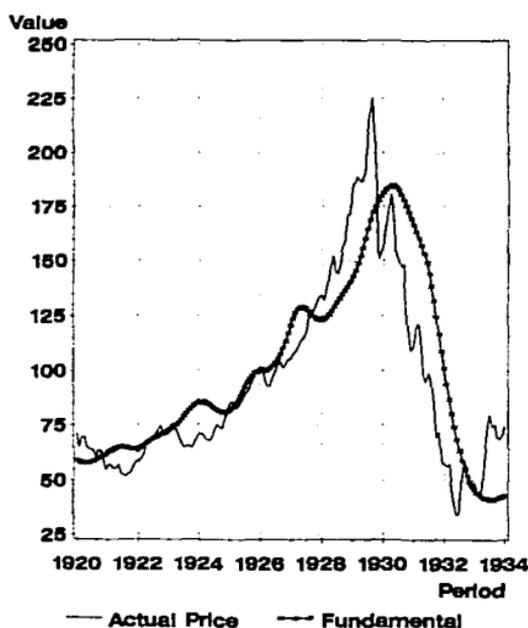


Figure 4
Actual price versus augmented Gordon fundamental

This figure plots the S&P 500 stock price index (plain line) versus the augmented Gordon growth model fundamental price (line joining stars) obtained by substituting Equation (5) into Equation (4) for g_t . See Appendix 1 for data sources.

Figure 4 plots actual prices along with the augmented Gordon fundamental prices produced by inserting Equation (5) into Equation (4) using the monthly equivalents of Barsky and DeLong's λ and r assumptions.¹⁰ From visual inspection of Figure 4, an opponent of bubbles would observe that the estimated fundamentals reproduce fairly well the broad movement in 1920s stock prices. Unfortunately, the market price is still sufficiently different from the estimated "fundamental" that a proponent of the bubbles hypothesis could claim a bubble in the market price. For example, notice from column 5 in part A of Table 1 that we fail to reject the null of a unit root in the augmented Gordon fundamental price error, thereby suggesting a persistent deviation of actual prices from fundamentals. Thus, assuming that the model in Equations (4) and (5) accurately reproduces the market's true fundamentals, we again find some support for the

¹⁰ In Figure 4, $\lambda = 0.9975$, which is the monthly equivalent of Barsky and DeLong's (1993) annual $\lambda = 0.97$, and $r = 0.0049$, the monthly equivalent of Barsky and DeLong's 6 percent annually. As in Barsky and DeLong, we set g_t equal to the average value of g over the entire period for which we have data, which in our case is 1899 to 1934.

bubbles hypothesis. However, we also reiterate our previous objection to the conclusion of a bubble in market prices on the grounds that, like previous models, the augmented Gordon model in Equations (4) and (5) is not sufficiently well specified to accurately capture all of the many subtleties in dividend behavior required to examine complicated phenomena such as price bubbles.

Misspecification in Barsky and DeLong's (1993) augmentation of the Gordon model comes from at least two sources. First, the Barsky and DeLong procedure uses assumed instead of estimated values for λ and r . Our attempts to estimate λ directly from the data reveal that while theory requires λ to be a number only slightly less than unity, maximum likelihood estimation of λ on dividend data from our entire time period of 1900 to 1933, as well as from the presample subperiod of 1900 to 1919, fails to yield a λ estimate even close to 0.97 annually (i.e., 0.9975 monthly). This suggests that, at least for the time period we study, the data do not support the imposed specification for g in Equation (5).¹¹ More important than this, however, is the misspecification error one produces by inserting Equation (5) into Equation (4) to obtain the augmented Gordon model. If g is nonconstant and evolves according to Equation (5), then Equation (4) cannot be true since, to derive Equation (4) from Equation (3), g is assumed constant. The augmented Gordon model comprised of Equations (4) and (5) is therefore internally inconsistent. This leads us to conclude that the augmented Gordon fundamentals do not match actual prices in Figure 4 because the model used to estimate fundamental prices is misspecified and not necessarily because there is a bubble in the market price.

2. A New Fundamentals Approximation Procedure

Although all three of the preceding traditional fundamentals estimation models are misspecified and give the appearance of bubbles, a definite progression in their relative ability to reproduce market data is evident. For example, the Gordon fundamental accounts for more of the rise and fall in market prices than does the expost fundamental because the Gordon model accounts for the nonstationarity of

¹¹ Recall that the assumed values used in Figure 4 are $\lambda = 0.97$ (annually), $r = 0.06$ (annually), and $g_0 = \bar{g}$ (which we calculate over 1899 to 1934). Using values of λ even a little bit different from the assumed value (e.g., using estimated values for λ) would result in a substantially less impressive fundamental series. Furthermore, using the constant discount rate of 8.3 percent obtained from our data (see Appendix 1), instead of Barsky and DeLong's $r = 0.06$, would produce a much flatter augmented Gordon fundamental that looks more like the basic Gordon fundamental in Figure 3. The augmented Gordon fundamental in Figure 4 is also rather sensitive to the assumed value of g_0 . In particular, we would obtain a significantly less impressive augmented Gordon fundamental if we restricted ourselves to in-sample information only when calculating g_0 .

dividends, while the expost model does not. Similarly, Barsky and DeLong's (1993) augmentation of the Gordon model performs even better than Gordon's original version because Barsky and DeLong allow the dividend multiplier, $(1 + g)/(r - g)$, to rise and fall with the dividend growth rate while in Gordon's original version the multiplier is constant. Thus, since both the level *and* growth rate of dividends rose during the late 1920s, the augmented Gordon model produces fundamentals that rise and fall more than the basic Gordon fundamental.

Our new approach rests on three observations taken from the three preceding models. First, we need to model growth rates, not levels. Second, these growth rates should be allowed to change over time, as they do in the actual data. Third, because growth is not constant, one cannot work with Gordon's Equation (4) but must instead use the less restrictive Equation (3) directly.¹² Given these observations, our objective is to find a well-specified model with which to estimate the market's expected value of the infinite sum of the progressive product of discounted dividend growth rates (i.e., $E_{t-1}\{y_t + y_t y_{t+1} + y_t y_{t+1} y_{t+2} + \dots\}$) so that this multiplier can be used directly with D_{t-1} in Equation (3) to obtain the fundamental price P_t^F .

2.1 The forecasting information set

To obtain our estimate of market's fundamental price, we will use only information available to time t investors—that is, information dated $t - 1$ and earlier—to (a) forecast the discounted dividend growth rate $y_t \equiv (1 + g_t)/(1 + r_t)$ into the infinite future and to then (b) calculate the progressive product of the forecasted y 's to obtain the date t multiplier and resulting fundamental price. Note that since our ultimate objective is to test for price bubbles, we cannot use past market prices to forecast future values for y . This is because, if there is a bubble in the market price, then by including past prices in our forecasting information set we might inadvertently impute a bubble into our fundamental price so that, on comparing our fundamental to the market price, we might erroneously fail to uncover the market price bubble. For this reason, we only use information on past discounted dividend growth rates, y_{t-i} , to forecast future discounted dividend growth rates, y_{t+j} .

In forming our discounted dividend growth series $\{y\}$ we can obtain y_t 's numerator, $(1 + g_t)$, directly from the data. However, we must make some assumptions regarding the investor's discount rate, $(1 + r_t)$, which forms the denominator of y_t . Cochrane (1992:252–253)

¹² See footnote 8 for an interesting special case in which Equation (4) can still be used.

suggests three interesting cases for $(1 + r_t)$: a constant discount rate, a consumption-based discount rate, and a discount rate equal to some variable reference return plus a risk premium. For completeness we calculate $(1 + r_t)$ in all three ways, each of which is explained more fully in Appendix 1.

For our constant discount rate we follow Shiller's (1989) approach and set r_t equal to the average real return on stocks for 1871 to 1988 (i.e., 8.13 percent annually, or 0.0067 monthly). For our consumption-based rate we follow authors such as Grossman and Shiller (1981) and employ the standard power utility function in consumption to obtain $(1 + r_t) = \beta^{-1}(C_t/C_{t-1})^\alpha$ where α is the coefficient of relative risk aversion, β is the subjective time discount rate, and C_t is consumption for the upcoming period t . Since most studies suggest that α somewhere between 1 and 2 is appropriate we use $\alpha = 1.5$, although our results are not overly sensitive to other reasonable values for α . We then set $\beta = 0.9953$ so that the average value of the consumption-based y s equals the average value of the constant- r y s. This gives an annualized $\beta = 0.945$, which is well within the range of $\beta \in [.90, .99]$ employed in most studies.¹³

For our reference-return plus risk-premium discount rate we set r_t equal to the real yield on high-quality short-term bonds in period t plus a constant equity premium equal to the average real return on stocks over high quality short-term bonds from 1871 to 1988. This particular approach has the attractive feature that $y_t = \frac{(1+g_t)}{(1+r_t)} = \frac{1+((G_t-\pi)/(1+\pi))}{1+((R_t-\pi)/(1+\pi))} = \frac{(1+G_t)}{(1+R_t)}$, where π is inflation and G_t and R_t are nominal dividend growth and discount rates, respectively. Thus, with r_t based on a time-varying bond rate, modeling y allows us to work with ratios of nominal discount and growth rates directly and thus removes a possible source of measurement error in the price index used to form π . A variety of statistical tests confirm that, for each of our three discounting conventions, the $\{y_t\}$ process is stationary over the presample interval of 1900 to 1919 as well as over the entire 1900 to 1933 period.¹⁴

2.2 The model

Given that y is stationary, the first stage of our fundamentals estimation procedure progresses in three standard steps: model specifica-

¹³ See Mehra and Prescott (1985) for some discussion of traditionally appropriate α and β values.

¹⁴ Phillips-Perron tests for a unit root in y , with data-dependent lag length, yield t -statistics of -2.58, -2.58, and -2.71 for constant- r , consumption- r and bond yield- r , respectively, over 1900 to 1919. For our entire sample of 1900 to 1933 we obtain t -statistics of -3.03, -3.06, and -2.94 for constant- r , consumption- r , and bond yield- r , respectively. In all cases the 10 percent critical value is -2.57. Augmented Dicky-Fuller tests yield even stronger rejections.

tion, parameter estimation, and forecasting. To specify the appropriate time-series model for y , we begin by analyzing the y data from before 1920, since this is presumably what an investor living in 1920 would have done in an effort to forecast y , and thus estimate fundamental prices into the late 1920s. We do not use data from after 1920 for the purpose of model selection since we do not want to use any information that was not available to market participants in the early 1920s, who were themselves trying to determine the correct model with which to form their own dividend forecasts.

Traditional Box-Jenkins analysis of y from 1900 to 1919 reveals that we need not search beyond ARMA models of order (2,2) to remove evidence of significant autocorrelation in y . However, LM tests reveal that there remains some significant autoregressive conditional heteroskedasticity in the model residuals. Furthermore, investigations for neglected nonlinearity of the type studied by Lee, White, and Granger (1993) also suggest the possibility of important nonlinear effects in the y processes. We found that, when these nonlinear effects are accounted for, nonlinearly augmented AR models with AR residuals fit the 1900 to 1919 data better than do simple ARMA models. To account for all the aforementioned features of the y process, and thus produce a correctly specified model, we therefore consider models for y that satisfy Equations (6) through (13) listed in Table 2.

The first summation in Table 2's Equation (6) contains the standard AR component of the y process. The second summation in Equation (6) is designed to capture nonlinear effects. As can be seen from Equation (7), the $\Psi(\cdot)$ terms take a logistic transform of the standardized differences (i.e., z) and squared differences (i.e., z^2) of y_{t-j} from its lagged value as specified in Equation (8). These logistic terms in the second summation in Equation (6) therefore mitigate, in a nonlinear manner, the persistence effects of outlier y s from the first summation in Equation (6). The ability of logistic terms like Equation (7)'s $\Psi(\cdot)$ to capture nonlinearities in a variety of applications has been well documented in the artificial neural network literature where their use has greatly proliferated in recent years.¹⁵ Finally, the residual from Equation (6) is modeled with the AR process in Equation (10) and the generalized ARCH conditional variance in Equations (11) and (12). A detailed discussion of the contribution and importance of each feature

¹⁵ See, for example, Kuan and White (1994). To achieve identification of the δ_i parameters in Equation (7), we follow convention in the artificial neural network literature and assign values to ω_k with Equation (9) so that ω_k fills the interval $[-1, 1]$ as k increases. Similarly, we multiply the parenthesized term in Equation (7) by 0.01 to satisfy the bound on the growth of the influence of the $\Psi(\cdot)$ term, as explained by Stinchcombe and White (1994). Work by Hornik, Stinchcombe, and White (1989, 1990) reveals that such a specification for $\Psi(\cdot)$ will allow the array of logistic terms to provide a universal approximator for a wide class of nonlinear functions.

Table 2
The ARAR-ARCH-ANN forecasting model

A. Definitions

- $D_t \equiv$ dividends paid during period t
- $g_t \equiv (D_t - D_{t-1})/D_{t-1} \equiv$ the dividend growth rate
- $(1 + r_t) \equiv$ the gross discount rate (constant, consumption-based, or bond yield-based)
- $y_t \equiv (1 + g_t)/(1 + r_t) \equiv$ the discounted dividend growth rate

B. The forecasting model

$$y_t = \alpha + \sum_{i=1}^{n_1} \beta_i y_{t-i} + \sum_{i=1}^{n_2} \Psi(\delta_i, z_t, \omega_t) + \epsilon_t \tag{6}$$

$$\Psi(\delta_i, z_t, \omega_t) = .01 \times \left(1 + \exp \left[\delta_i \left(\omega_{\Delta y} + \sum_{j=1}^l \{ \omega_{1ij} z_{t-j} + \omega_{2ij} z_{t-j}^2 \} \right) \right] \right)^{-1} \tag{7}$$

$$z_{t-j} = (y_{t-j} - y_{t-j-1})/\sigma_{\Delta y} \tag{8}$$

$$\omega_{\Delta y} = \sin(\pi \times \ln |s + k^2 + i + j|) \tag{9}$$

$$\epsilon_t = \sum_{i=1}^q \rho_i \epsilon_{t-i} + u_t \tag{10}$$

$$u_t = \sqrt{h_t} \eta_t \quad : \quad \eta_t \sim (0, 1) \tag{11}$$

$$h_t = \lambda + \sum_{j=1}^{m_1} \phi_j h_{t-j} + \sum_{j=1}^{m_2} \xi_j u_{t-j}^2 \tag{12}$$

$$n_1, n_2, s, l, q, m_1, m_2 \in \{0, 1, 2\} \tag{13}$$

This table presents the model used in our forecast simulation exercises outlined in the text and in Appendix 2. Values of $n_1, n_2, s, l, q, m_1,$ and m_2 for the constant- r , consumption- r and bond yield- r y series are optimally chosen on 1900 to 1919 data using the augmented Box-Jenkins procedure outlined in Appendix 2. The optimal specifications are as follows: constant- r : $n_1 = 2, n_2 = 1, s = 1, l = 2, q = 1, m_1 = 0, m_2 = 1$; consumption- r : $n_1 = 2, n_2 = 1, s = 2, l = 2, q = 1, m_1 = 0, m_2 = 1$; bond yield- r : $n_1 = 1, n_2 = 1, s = 1, l = 1, q = 2, m_1 = 0, m_2 = 1$.

in our model for accurately forecasting discounted dividend growth is presented in Section 3.

Equation (13) provides specification grid boundaries for the models that we consider. Specific values of $n_1, n_2, s, l, q, m_1, m_2$ for the constant- r , consumption- r , and bond yield- r y series are optimally chosen on 1900 to 1919 data using the augmented Box-Jenkins procedure outlined in Appendix 2. The optimal specifications are as follows: constant- r : $n_1 = 2, n_2 = 1, s = 1, l = 2, q = 1, m_1 = 0, m_2 = 1$; consumption- r : $n_1 = 2, n_2 = 1, s = 2, l = 2, q = 1, m_1 = 0, m_2 = 1$; bond yield- r : $n_1 = 1, n_2 = 1, s = 1, l = 1, q = 2, m_1 = 0, m_2 = 1$. A variety of tests reveal that these models are all well specified, as one would expect from our model construction procedure (see Appendix 2 for details). The robustness of our results to a variety of specification

errors caused by omission of various terms in Equations (6) through (13) is discussed below.

Holding fixed the model specifications chosen with 1900 to 1919 data, we next proceed to parameter estimation. We use a rolling estimation procedure to produce, for each date from 1920 to 1933, parameter estimates conditional only on information available to the investor at the time prices were being determined. For example, parameter values used to forecast future discounted dividend growth as of 1920:1, and thus calculate the fundamental price for 1920:1, are estimated using only data on y up to 1919:12. Parameter values used to calculate the fundamental price for 1920:2 are then estimated using only data up to 1920:1, and so forth. We never use data from after $t - 1$ to obtain parameter estimates at time t since this would give our model an unfair advantage over investors who were actually determining market prices during the 1920s. We do update our parameter estimates at each date as new information becomes available, however, since this new information would undoubtedly have been used by investors in an effort to increase the accuracy of their estimates of the economy's "true" underlying parameters.

Finally, we use the model specification and parameter estimates, which were formed using only information available to the market at time t (i.e., information dated $t - 1$ and earlier) to forecast the sequence y_t, y_{t+1}, y_{t+2} , etc. Our goal in doing this is to calculate, for each date t , the expected value of the sum of the progressive product of the forecasted y 's that appears on the right-hand side of Equation (3), and then to multiply known dividends from date $t - 1$ by this sum to obtain the estimated fundamental P_t^F .

As can be seen from the model in Equations (6) through (13), y_{t+i} is not generally independent of y_{t+j} so that $E_{t-1}\{y_t y_{t+1}\} \neq E_{t-1}\{y_t\} E_{t-1}\{y_{t+1}\}$ in general. Thus, when calculating expectations of the bracketed term on the right side of Equation (3), we cannot simply set $u_{t+i} = 0$ in Equation (10) and extend out the y series in the traditional forecasting fashion to form a separate expectation for each individual y_{t+i} . Instead, the expected value of the sum of the progressive product in Equation (3) must be calculated numerically using Monte Carlo simulation. This involves producing cross sections of time series (up to 10,000 observations long) for u in Equation (10), which are randomly drawn from the distribution in Equations (11) and (12), to produce cross sections of simulated time series for ϵ in Equation (10) and thus y in Equation (6). Each simulated time series for y is then progressively multiplied and summed as in Equation (3). This procedure is repeated 10,000 times to produce 10,000 different values for the bracketed term on the right-hand side of Equation (3). The expected value of the bracketed term is then calculated as the mean

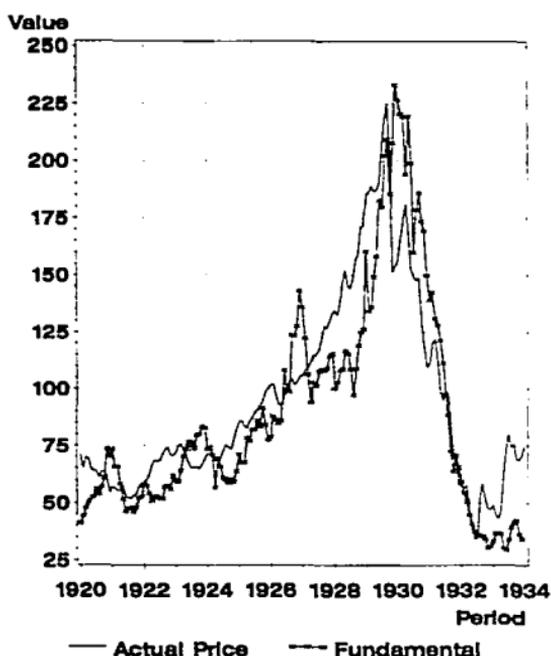


Figure 5
Actual price versus forecast simulation price (constant- r)

This figure plots the S&P 500 stock price index (plain line) versus the fundamental price (line joining stars) obtained with our ARAR-ARCH-ANN model in Table 2 and forecast simulation procedure from Appendix 2. The discount rate r_t is a constant equal to the average return on stocks from 1871 to 1988. See Appendix 1 for data sources.

of the 10,000 simulated summed products. This expected value is finally multiplied by the most recently observed dividend to calculate the fundamental price P_t^F . A detailed description of our simulation procedure is contained in Appendix 2.

2.3 Results

The solid line in Figure 5 plots the actual S&P 500. The line connecting stars in Figure 5 plots our forecast simulation fundamental price sequence produced with constant discount rates. Note that our forecast simulation fundamental price rises and falls like actual prices and peaks within 10 points of the actual price's peak. Although our fundamental peaks 3 months later than actual prices (i.e., in December 1929, instead of September 1929), our fundamental clearly rejects the notion of an explosive bubble that expands until September 1929 and then bursts. An explosive bubble requires that the difference between the market and fundamental price increase as long as the bubble floats. Conversely, Figure 5 reveals that the difference between our fundamental and the market price is actually shrinking throughout 1928 and

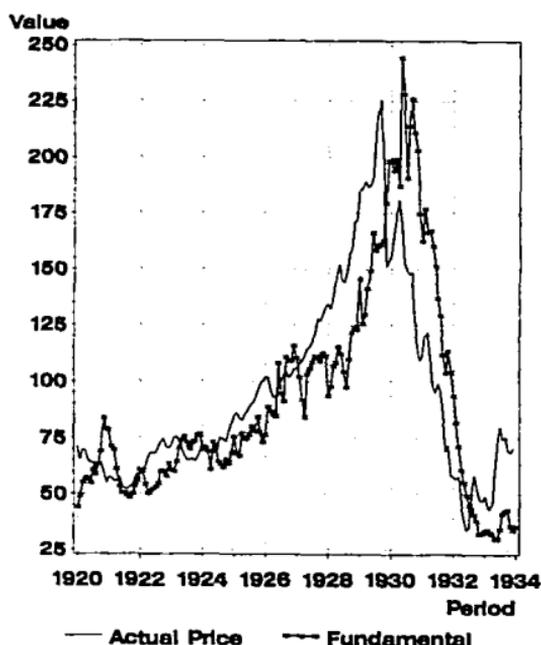


Figure 6
Actual price versus forecast simulation price (consumption- r)

This figure plots the S&P 500 stock price index (plain line) versus the fundamental price (line joining stars) obtained with our ARAR-ARCH-ANN model in Table 2 and forecast simulation procedure from Appendix 2. The gross discount rate $(1 + r_t)$ is derived from a standard power utility function in consumption with the coefficient of relative risk aversion $\alpha = 1.5$ and the monthly subjective time discount rate $\beta = 0.9953$ (i.e., 0.945 annually). See Appendix 1 for data sources.

1929. Indeed, the Phillips-Perron unit root t -statistic in the last row of part A in Table 1 (column 6) reports that, unlike the traditional models, we are able to reject the null hypothesis that our real price error displays a unit root (the same result is obtained if we restrict our time period to 1920 to 1929). This suggests that there is *not* a bubble in actual prices when compared to our fundamentals. Furthermore, as theory suggests should be the case in the absence of bubbles [e.g., Camerer (1989), Shiller (1989)], part B of Table 1 reveals that returns from our fundamental series are slightly more volatile than actual returns, as measured by their standard deviation, and have fourth moments (i.e., tail thickness) roughly equal to actual returns. In addition, our fundamental implies returns that, like actual returns, are stationary. Thus, unlike the traditional models we investigate, our dividend forecasting model produces returns with properties similar to actual returns.

Figure 6 plots actual prices along with our consumption-based forecast simulation fundamental price. Although the consumption-based

fundamental does rise and fall like actual prices, this fundamental does not appear to mimic actual prices quite as well as the constant interest rate fundamental. In particular, we see from the last row of part A in Table 1 (column 7) that, unlike the constant- r case, the consumption- r fundamental fails to reject a unit root in the real price error at the 10 percent level. However, our failure in this one test is weak (estimated t -ratio = -2.35, 10 percent critical value = -2.57) and we see no strong evidence of bubbles in any of our other statistical examinations of the consumption- r fundamental. For example, from part B of Table 1 (column 7) we see that the moments of the fundamental returns distribution accord well with the moments of actual returns. The observation that our consumption- r fundamental has variance and tail thickness roughly equal to actual returns is especially interesting given the well-documented difficulty of traditional procedures to reconcile the volatility of asset prices with the smoothness of consumption. For example, while the results from our new forecasting model are obtained with a risk aversion parameter of $\alpha = 1.5$, an α value well within the range of 1 to 2 suggested by most theoretical studies, Grossman and Shiller's (1981) traditional approach achieves only marginal success with $\alpha = 4$, and Mehra and Prescott (1985) find that even with $\alpha = 10$ it is difficult to reconcile returns volatility with consumption smoothness.

Figure 7 plots actual prices along with our forecast simulation fundamental price with a discount rate equal to the real high-quality short-term bond rate plus a constant risk premium. The similarity between the actual price series and this fundamental series is truly remarkable. The dividend forecasts from our model, which are based only on dividend and interest rate information available to investors at the time actual prices were being determined, produce a fundamental price series that rises from roughly the same value as actual prices in 1920 to peak at exactly the same time, and within 10 points of, the actual price in 1929. Our fundamental price then crashes along with the observed price in the early 1930s and recovers with actual prices in the mid 1930s. There are several reasons for this interesting result, all of which we discuss in the following section. However, before doing so, it is useful to confirm that Table 1's statistical results support our visual findings. From column 8 of Table 1 we see that Figure 7's real price error is indeed stationary and that the first four moments of the fundamental returns distribution match those from the actual market returns. In particular, fundamental returns are slightly more volatile than actual returns and almost as fat-tailed. In short, there is no evidence of a bubble.

An investor living in the 1920s would undoubtedly have used a wide variety of information (e.g., business forecasts, company reports,

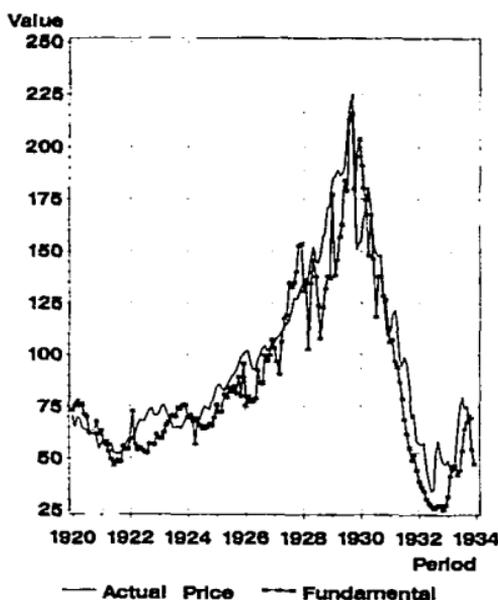


Figure 7

Actual price versus forecast simulation price (bond yield- r)

This figure plots the S&P 500 stock price index (plain line) versus the fundamental price (line joining stars) obtained with our ARAR-ARCH-ANN model in Table 2 and forecast simulation procedure from Appendix 2. The discount rate r_t equals the interest rate on low-risk commercial paper plus a constant equity premium equal to the average excess return on stocks over low-risk bonds from 1871 to 1988. See Appendix 1 for data sources.

past prices, etc.) to form expectations of future dividends. We have only used a small subset of this information: dividends alone in Figure 5, dividends and consumption in Figure 6, and dividends and interest rates in Figure 7. We would therefore not expect our fundamentals to exactly match the true fundamentals constructed by 1920s' investors.¹⁶ Thus, even if there was no bubble in actual prices, we would not expect our fundamentals to exactly match actual market prices. Nevertheless, from visual examination of Figures 5 through 7 and the statistics in Table 1, it is difficult to claim that market prices differ from our estimates of fundamental prices in a substantial and persistent bubble-like manner. Since the models used to obtain our fundamental prices are well specified and derived from the dividend data, and are therefore consistent with this data, we therefore reject the hypothesis that there was a bubble in the 1920s stock market.

¹⁶ Indeed, one might argue that our estimated fundamental prices should lag behind actual prices to the extent that our models require extra y realizations to compensate for the fact that our forecasts are based on a rather limited subset of all information available to investors.

3. Discussion

Two observations arise from Table 1 and the preceding figures. First, our new forecasting procedure yields fundamentals that reject bubbles, while traditional models do not. Second, on comparing Figures 5, 6, and 7 we see that, among our three forecasting models, the bond yield discount rate specification for y performs best and the consumption-based y worst. In this section we attempt to explain these findings.

3.1 Modeling freedom and dividend forecasts

One could argue that our model is more successful than traditional models in mimicking market prices because our new dividend forecasting procedure more accurately reproduces the market's true expectations of future dividends. Unfortunately, such an assertion can never be tested directly since we cannot observe the market's true dividend expectations. It is nevertheless useful to begin our discussion by examining the ability of each model to forecast dividends from some date t into the infinite future, given only information from date $t - 1$ and earlier. We then compare the forecasts from each model with realized dividends. While in principle any date could be chosen as the forecast starting point, we begin with an example in which date t is June 1932, the month in which market prices reached their lowest point following the crash. This choice is especially useful for illustrative purposes because it allows us to examine the various models' abilities to predict the trough and turning point in actual dividends, which occurred in mid-1933 (see Figure 3).

Figure 8 plots 19 months of realized dividends from May 1932 (period 0) to December 1934, the last month in our data sample, along with the first 60 months of out-of-sample dividend forecasts made by various models; that is, forecasts for periods 1 through 60 conditional only on information from period 0. Dividend forecasts for each model are therefore obtained as $D_{t+i} = D_{t-1} \prod_{j=0}^i (1 + g_{t+j})$, where $t-1 =$ May 1932 and g_{t+j} is the period $t + j$ rate of dividend growth forecasted by the model in question.

The plain curve in Figure 8—which ends in period 18: December 1934—plots actual dividends. This curve reveals that actual dividends declined from June 1932 (period 1) until June 1933 (period 13), after which they began to rebound. This realized pattern looks very different from the monotonically increasing top two curves in Figure 8, which plot dividend forecasts obtained from the augmented Gordon and basic Gordon models (top and second-from-top curves, respectively). Indeed, compared to actual dividends, the Gordon models are unrealistically bullish, forecasting increasing dividends from June 1932

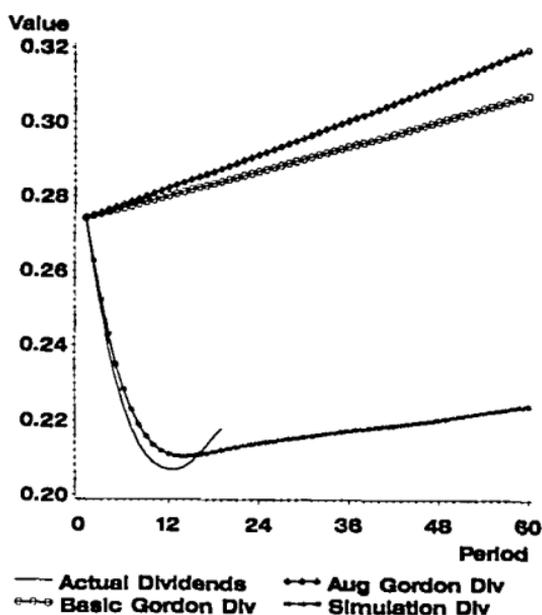


Figure 8
Dividends forecasted as of date 1 = 1932: 6

This figure plots the sequence of future dividends forecasted by various models as of 1932:6 (period 1) using only information from 1932:5 (period 0) and earlier. The truncated plain line plots actual dividends. The line joining circles plots dividends forecasted by the basic Gordon growth model of Equation (4). The line joining diamonds plots dividends forecasted by the augmented Gordon growth model, obtained by inserting Equation (5) into Equation (4). The line joining stars plots dividends forecasted by our ARAR-ARCH-ANN forecast simulation procedure, outlined in Appendix 2, using the model in Table 2 with a constant discount rate r_t equal to the average return on stocks from 1871 to 1988. See Appendix 1 for data sources.

onward. Quite a different result is seen in the second-from-bottom curve (joining stars) in Figure 8, which plots the sequence of dividends forecasted by our new procedure with a constant discount rate as in Figure 6.¹⁷ This forecast simulation dividend curve reveals that, using only dividend information available in May 1932, our new model predicts that dividends will fall for the next 14 months and then slowly rise thereafter, a pattern very similar to that subsequently taken by realized dividends.

The ability of our model to produce out-of-sample dividend forecasts that fall and then rise like realized dividends highlights an im-

¹⁷ See Section 2 and Appendix 2 for details on exactly how our forecasts are produced. Note that what is actually plotted in Figure 8 is the mean of the dividend forecasts for each date from 10,000 replications of our dividend forecasting model. Also note that we use the constant- r version of our model here because our procedure forecasts $y_{t+j} = (1 + g_{t+j}) / (1 + r_{t+j})$, and thus we cannot in general separate forecasts of dividend growth from forecasts of the discount rate. However, with $(1 + r_{t+j})$ constant as in Figure 6, we can easily retrieve the forecasted numerator, $\prod_{j=0}^t (1 + g_{t+j})$, and thus a dividend forecast for each period from date t into the infinite future.

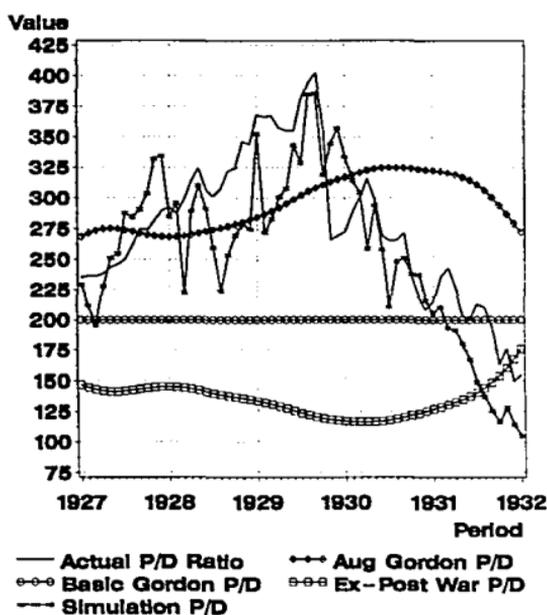


Figure 9
Monthly price-dividend ratios

This figure plots *monthly* price-dividend (P/D) ratios from various models. The plain line (highest line at the 1929 peak) plots the actual P/D ratio. The line joining stars (second to highest line at the 1929 peak) plots the fundamental P/D ratio where prices are obtained with our ARAR-ARCH-ANN forecast simulation procedure, outlined in Appendix 2, using the model in Table 2 with the discount rate r_t equal to the interest rate on low-risk commercial paper plus a constant equity premium equal to the average excess return on stocks over low-risk bonds from 1871 to 1988. The line joining diamonds (third to highest line at the 1929 peak) plots the augmented Gordon fundamental P/D ratio obtained by substituting Equation (5) into Equation (4). The line joining circles (horizontal line) plots the basic Gordon fundamental monthly P/D ratio which, from Equation (4), is a constant at $(1 + g)/(r - g) = 200$. The line joining squares (bottom line) plots the ratio of ex post warranted prices to dividends.

portant difference between our procedure and the familiar Gordon model. When using Gordon's model from Equation (4) to forecast future cash flows, one must assume that dividends will grow at the same constant rate into the infinite future. This is true even if one allows the assumed constant rate to change each time the forecasting procedure is employed—that is, for each new starting date T —as in Barsky and DeLong (1993). Conversely, since our approach works directly with Equation (3), we do not need to make Gordon's constant growth assumption. Our new procedure therefore forecasts future growth rates in a manner that allows for the possibility of a different forecasted growth rate in each future period. Figure 8 reveals that this extra degree of freedom in modeling y 's behavior helps us produce superior forecasting results.

The extra modeling freedom offered by our procedure is especially important when calculating the expected discounted present value of

the sequence of forecasted dividends; that is, $E_{t-1} \left\{ \sum_{k=0}^{\infty} (\Pi_{i=0}^k y_{t+i}) \right\}$ on the right side of Equation (3). Indeed, by rewriting Equation (3) to note that $E_{t-1} \{ \cdot \} = P_t^F / D_{t-1}$, it is easily seen that the mechanical explanation for why our new model outperforms traditional models is that our dividend multiplier, $E_{t-1} \{ \cdot \}$, looks more like the market's realized price-dividend ratio than do the multipliers implied by traditional fundamentals models. To see this, Figure 9 plots the monthly price-dividend (P/D) ratios (or multipliers) for the fundamentals from Figures 2, 3, 4, and 7 (note that these *monthly* P/D ratios will be an order of magnitude greater than traditionally reported *annual* P/D ratios). Correlations between Figure 9's actual P/D ratio and the other fundamental P/D ratios, or dividend multipliers, are reported in Table 1.

The actual P/D ratio is the plain line in Figure 9 which peaks in September 1929. The horizontal line connecting circles at 200 in Figure 9 is the constant basic Gordon monthly multiplier, $(1+g)/(r-g) = 200$ from Equation (4), as explained in Section 1.3. The bottom curve connecting squares, with a trough instead of peak in 1930, is the ratio of Figure 2's ex post warranted prices to dividends. The negative correlation between the ex post and realized P/D ratios is to be expected since, on comparing Figures 2 and 3, we see that the ex post warranted price is even flatter than dividends. The smoothly rising and falling curve connecting diamonds that peaks a full year after, and considerably lower than, the actual price-dividend peak is the augmented Gordon multiplier. The augmented Gordon multiplier's late peak, which is a product of Equation (5)'s slow-to-react long-lag weighting mechanism, explains why Figure 4's augmented Gordon fundamental peaks later and lower than the market price. Finally, the jagged line connecting stars, which most closely resembles the actual P/D curve, is the $E_{t-1} \{ \cdot \}$ multiplier—that is, the fundamental P/D ratio—from our model.¹⁸ From Figure 9 and the statistics in Table 1 it is easily seen that our model produces fundamental prices that behave most like market prices because our model produces a time series of dividend multipliers, $E_{t-1} \{ \cdot \}$, that most closely mimics the market P/D ratio.

3.2 Model specification

So far we have seen that the freedom to model directly the time-series properties of discounted dividend growth, y , is a key factor for forecasting future dividends, and thus for calculating the expected discounted present value of the sequence of forecasted future dividends,

¹⁸ As one can surmise from visual inspection of Figures 5 and 6, multipliers from our constant- r and consumption- r models look substantially similar to those from the bond yield- r shown in Figure 9, except that the constant- r and consumption- r multipliers peak a little later.

$E_{t-1} \left\{ \sum_{k=0}^{\infty} \left(\prod_{i=0}^k \gamma_{t+i} \right) \right\}$. However, the increased freedom afforded by our direct use of Equation (3), instead of the restrictive Equation (4), can only be taken advantage of if we employ a well-specified model to forecast the y sequence. This fact can be clearly demonstrated by observing the effects of omitting various features of Equations (6) through (13) and instead using a misspecified model to estimate fundamental prices. For the sake of brevity, we will concentrate on the bond yield- r version of our model, whose multiplier $E_{t-1}\{\cdot\}$ is plotted in Figure 9. Results from the constant- r and consumption- r versions are substantially similar.

3.2.1 The basic ARAR shell. The most basic strategy one could employ when modeling the y_t series is to simply regress y on a constant, so that the forecasted future y_{t+i} s are just an equally weighted moving average of the past y 's. Not surprisingly, the fundamental prices we obtain from such a model do not track market prices at all well. Indeed, the long-run average value of y_t changes so slowly that the fundamental prices from this overly simple model (not plotted in this article) look very similar to the basic Gordon fundamental in Figure 2.

A somewhat more sophisticated modeling approach is to employ a standard time-series representation for y_t in an effort to capture more of y 's time-series properties. We therefore plot in Figure 10 the fundamental price series obtained with dividend forecasts from the basic ARAR shell of a model that remains after we omit from Equations (6) through (13) the nonlinear ANN (artificial neural network) terms in Equations (7) through (9) and the ARCH term in Equation (12).¹⁹ In other words, we plot in Figure 10 the bond yield- r fundamentals from the forecast simulation model given by $y_t = \alpha + \beta_1 y_{t-1} + \epsilon_t$; $\epsilon_t = \sum_{i=1}^2 \rho_i \epsilon_{t-i} + u_t$; $u_t \sim (0, \sigma)$.

As one would expect, the ARAR fundamental price series in Figure 10 rises and falls somewhat more than the fundamental price series from Figure 2's constant y_t Gordon model. However, even a carefully chosen ARAR model is still sufficiently misspecified that the fundamental price series it produces peaks 75 points lower than actual market prices in September 1929. Indeed, if we eliminate from our full model in Equations (6) through (13) both the ARCH residual and

¹⁹ Recall from our discussion in Section 2.2 that we employ an AR in mean with AR residuals as our base model for y_t (see Equations (6) through (13) in Table 2), instead of ARMA, because with the logistic ANN terms and ARCH included the ARAR base representation fits the data better than ARMA. To facilitate easy comparison with our article's other figures and results, which are based on the ARAR foundation, we therefore plot in Figure 10 the fundamental prices produced by the ARAR shell of our model which excludes both ARCH and ANN terms. Results from a basic ARMA specification are essentially the same as those reported for the ARAR specification.

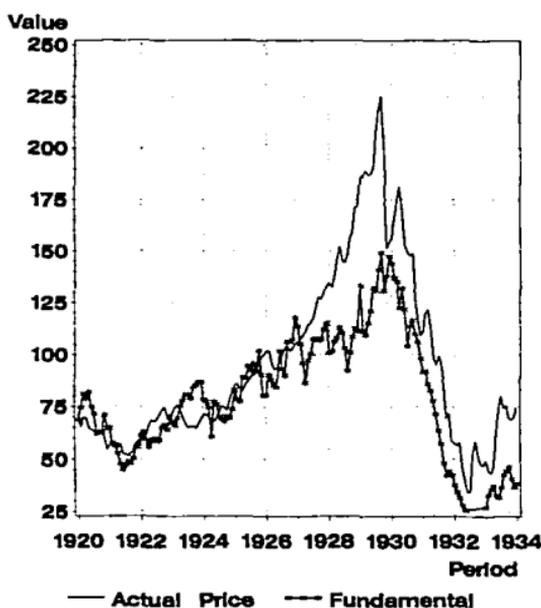


Figure 10
Actual price versus forecast simulation price (bond yield- r): ARAR shell

This figure plots the S&P 500 stock price index (plain line) versus the fundamental price (line joining stars) obtained with our forecast simulation procedure, outlined in Appendix 2, using the model in Table 2 with the ARCH and ANN terms excluded to leave only the basic ARAR shell remaining. The discount rate r_t is the interest rate on low-risk commercial paper plus a constant equity premium equal to the average excess return on stocks over low-risk bonds from 1871 to 1988. See Appendix 1 for data sources.

the nonlinear $\Psi(\cdot)$ term, then the rudimentary ARAR shell of a model that remains produces fundamental prices in Figure 10 that, for several of our tests, suggest a bubble in market prices. It is important to note, however, that any appearance of bubbles from a standard time-series model can be explained by the fact that even a carefully chosen $AR(p)AR(q)$ (or $ARMA(p,q)$) model for y_t fails specification tests for residual ARCH and nonlinear effects and can therefore not be expected to capture all the features of the data necessary to produce reliable discounted dividend forecasts.

3.2.2 The importance of ARCH. We now add the nonlinear logistic ANN $\Psi(\cdot)$ terms in Equations (7) through (9) back into our model, but still omit the ARCH terms in Equation (12); that is, we reestimate Table 2's model Equations (6) through (13) under the *false* assumption that the residual from Equation (10) is homoskedastic. The resulting no-ARCH fundamental price series is plotted in Figure 11. Note that Figure 11's no-ARCH fundamental performs better than the simple ARAR shell in Figure 10, but not nearly so well as Figure 7's fully

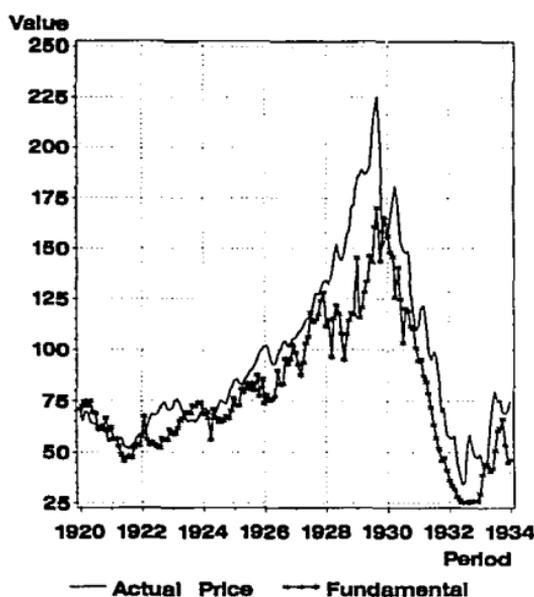


Figure 11
Actual price versus forecast simulation price (bond yield- r): no ARCH

This figure plots the S&P 500 stock price index (plain line) versus the fundamental price (line joining stars) obtained with our forecast simulation procedure, outlined in Appendix 2, using the model in Table 2 with the ARCH terms excluded. The discount rate r_t is the interest rate on low-risk commercial paper plus a constant equity premium equal to the average excess return on stocks over low-risk bonds from 1871 to 1988. See Appendix 1 for data sources.

specified model that includes ARCH. In particular, without ARCH in the model (Figure 11) we are unable to capture the full height of the 1929 price peak.

Mechanically, the model with ARCH included (Figure 7) does a better job of catching the full height of the 1929 peak than the no-ARCH restricted version (Figure 11) because the conditional variance of Equation (10)'s u_t residual, which is given by b_t from Equation (12), is falling over time. This fact is revealed in Figure 12, which plots at each date the cross-sectional average of the forecasted b_t variances produced by our 10,000 model simulations. Note in particular that the b_t variance is much lower in 1929 than it is in 1920 and is lower again by the mid-1930s.

To see why the ARCH model's falling innovation variance is important for forming fundamental prices, note that the variance of our y_t series can be decomposed into the sum of the variance of the autocorrelated component of y_t , which we model with Equations (6) through (10), plus the variance of Equation (10)'s u_t residuals; that is, $\text{Var}(y_t) = \text{Var}(y_t - u_t) + \text{Var}(u_t)$, where $\text{Var}(u_t) = b_t$ from Equation

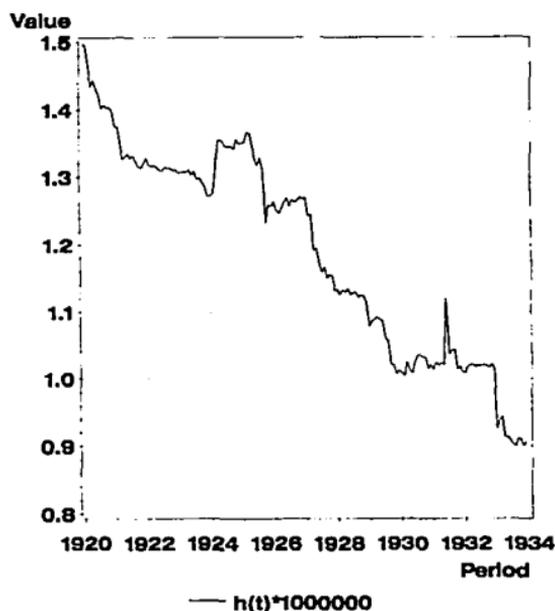


Figure 12
Forecasted error variance

This figure plots the average value of the b_t conditional variance from our 10,000 simulations at each date. Appendix 2 outlines the simulation procedure used. Table 2 contains the model employed. The discount rate r_t is the interest rate on low-risk commercial paper plus a constant equity premium equal to the average excess return on stocks over low-risk bonds from 1871 to 1988. See Appendix 1 for data sources.

(12). Stationarity of the y_t process ensures that in the limit the mean of our forecast-simulated y_t distribution converges to the unconditional mean of y_t and is therefore independent of the initial conditions in our simulation exercise, which are determined by the values of $y_{t-1}, y_{t-2}, \epsilon_{t-1}, \epsilon_{t-2}$ from Equations (6) through (13) as explained in Appendix 2. In particular, the value of $E_{t-1}\{y_{t+n}\}$ as $n \rightarrow \infty$ is asymptotically independent of the values of (y_{t-1}, y_{t-2}) on which our 10,000 simulated economies are based.

Note, however, that the speed with which the mean of our cross section of simulated y_{t+i} s converges to the unconditional mean of y_t depends on the variance of the autocorrelated component of y_t relative to the variance of the u residuals. Obviously, if the b_t residual variance completely dominated the autocorrelated component's variance, then the y_{t-i} initial conditions would be completely unimportant for forecasting future y s, as the mean of the y_{t+i} s generated at each date i by our 10,000 simulated economies would simply equal the unconditional mean of y_t . Conversely, with some importance being attached

to the autocorrelated component of y , the cross-sectional mean of the simulated y_{t+i} s will only gradually converge to the unconditional mean as $i \rightarrow \infty$.

The speed with which the cross-sectional mean of the simulated y_{t+i} s converges to the unconditional mean of y is inversely related to the magnitude of the residual's variance, b_t . Thus, as b_t declines and the importance of the u residual decreases relative to y 's autocorrelated component, the simulated series' speed of convergence to the unconditional mean of y declines. As b_t falls during the late 1920s, past realized values of y in the data—that is, the initial conditions in our simulations—therefore have an increasingly persistent influence on the simulated evolution of y_{t+i} . Thus, with falling b_t , initial conditions have a greater influence on the sum of the progressive product $E_{t-1}\{\cdot\}$ on the right side of Equation (3) and therefore on the fundamental price produced by our procedure.

Notice from the slope of the dividend series plotted in Figure 3 that the growth rate in dividends (i.e., slope) increased during the latter half of the 1920s (this fact will be demonstrated more fully below). Since b_t is assumed constant in the model without ARCH, the no-ARCH model interprets this late 1920s increased dividend growth in the same manner as it interpreted movements in y from the early 1920s when the innovation variance was high. The constant variance no-ARCH model therefore produces simulated future values for y_{t+i} that revert to the unconditional mean of y rather quickly. Thus, the no-ARCH model produces values for Equation (3)'s $E_{t-1}\{\cdot\}$ multiplier during the late 1920s that are somewhat similar to values produced by the no-ARCH model in other periods.

Conversely, in our full model with ARCH included, the late 1920's drop in b_t (see Figure 12) increases the influence of the y_{t-1} , y_{t-2} initial conditions on the evolution of the y_{t+i} simulations. Large realizations for y_{t-1} , y_{t-2} in the late 1920s therefore keep our ARCH forecasted y_{t+i} s above the unconditional mean longer than would be the case with no-ARCH constant variance. Combined with the late 1920's increased y_t s, the late 1920's smaller b_t ARCH variance therefore yields a larger $E_{t-1}\{\cdot\}$ multiplier than would be the case with constant variance. This is one reason (others are discussed below) why our full model with ARCH in Figure 7 produces a higher fundamental price than does our no-ARCH model in Figure 11.²⁰

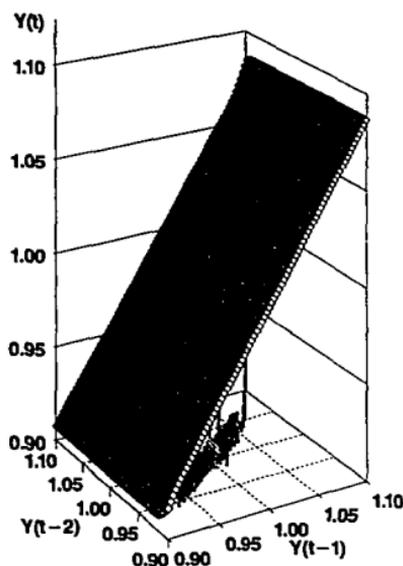
²⁰ In addition, the continuing decline in the variance of u_t innovations throughout the 1930s causes our full model with ARCH to react more strongly to the postcrash decline in dividend growth than does the constant variance no-ARCH model. This helps to explain why the fully specified fundamental in Figure 7 falls farther between 1929 and 1932 than does the constant variance fundamental in Figure 11.

3.2.3 The importance of ANN terms. An important problem with omitting ANN terms from our model in Equations (6) through (13) is that, without the nonlinear terms in Equations (7) through (9), the conditional mean process for y_t is misspecified. Thus, whatever parameters are included in the model will be inconsistently estimated and the model residual will contain systematic effects that would otherwise have been captured by the omitted variable. The conditional variance process for the u_t residual from Table 2's model without ANN terms will therefore embody two effects: a spurious variance effect produced by the conditional mean misspecification and a true variance effect. Because of the confounding spurious effect, it may be difficult for the ARCH terms in a no-ANN model to accurately capture the true conditional variance process necessary to obtain true fundamental prices. Thus, without ANN in the mean, it may be difficult for ARCH to perform its intended function described in the previous section. Studying the performance of our model without ANN terms nevertheless provides some interesting insight.

If we do omit the ANN terms in Equations (7) through (9) from our model, but include ARCH, then we obtain a no-ANN (but ARCH included) fundamental price series that looks very similar to the basic ARAR fundamental in Figure 10. Indeed, this similarity is so strong that we do not present a separate plot of the no-ANN (but ARCH included) fundamental price series in our article. This close similarity between the ARAR and no-ANN (but ARCH included) fundamentals reveals that, unless ANN terms are included in the model, there is almost no benefit to adding ARCH. It thus appears that a key role of the ANN terms is to correctly specify the conditional mean of y_t so that the ARCH variance process can be estimated correctly.

One way to visually observe why the ANN terms are so important for correctly specifying the conditional mean of y_t is to plot the forecasting surface implied by our model with ANN included. We therefore graph in Figure 13 the three-dimensional surface which is the bond yield- r model's forecast mapping from y_{t-1} , y_{t-2} into y_t , as specified in Equations (6) through (13) with parameter estimates obtained using data available to the market in September 1929. For tractability, and to isolate the ANN effect, we set the ϵ forecast error in Equation (6) to its unconditional mean of zero so that the forecasting surface plotted in Figure 13 is given by $y_t = 0.145 + 0.85y_{t-1} + 0.01(1 + \exp[3.41(-0.31 - 0.94[y_{t-1} - y_{t-2}]/0.02 - 0.17[y_{t-1} - y_{t-2}]^2/0.02^2]))^{-1}$.

From Figure 13 and the preceding equation, we see that the forecasting surface is comprised of three sections. First, the left side of Figure 13—in the region where $y_{t-1} < y_{t-2}$ (e.g., $y_{t-1} = 0.95$, $y_{t-2} = 1.05$)—is occupied by a lower triangle-shaped plane with its wide base along the bottom of the figure's y_{t-2} axis and its hypotenuse up



Forecasting Surface: Diamonds

Data Points: Pillars

Figure 13

Forecasting function: September 1929

This figure plots, as the surface of diamonds, the functional relationship between (y_{t-1}, y_{t-2}) and y_t implied by our model in Table 2, with the discount rate r_t equal to the interest rate on low-risk commercial paper plus a constant equity premium equal to the average excess return on stocks over low-risk bonds from 1871 to 1988. Model parameter values are from September 1929, the month of the stock market crash. The 60,000 short vertical pillars rising up from the figure's floor represent the first 6 months of Monte Carlo evolutions for y_{t+1} (measured on the vertical axis) as a function of y_{t+1-1}, y_{t+1-2} (i.e., the two most recently forecasted values of y in the Monte Carlo sequence) from our 10,000 simulated economies produced in September 1929. The five tall pillars that extend from the figure's floor all the way up through the forecasting surface represent, from top to bottom, the maximum, 99.5th percentile, median, 0.5th percentile, and minimum values of y_{t+1} forecasted in September 1929.

the diagonal of the forecasting surface from the lower front corner to the top back corner, where $y_{t-1} \approx y_{t-2}$. Second, on the right side of the figure—where $y_{t-1} > y_{t-2}$ (e.g., $y_{t-1} = 1.05, y_{t-2} = 0.95$)—is an upper triangle-shaped plane with its wide base along the top of the figure and its hypotenuse on the surface's diagonal, where $y_{t-1} \approx y_{t-2}$. Notice that this upper triangular plane is shifted higher up the y_t axis than is the lower triangle; that is, the value of y_t forecasted with $y_{t-1} = 1.05, y_{t-2} = 0.90$ is higher than the value of y_t forecasted with $y_{t-1} = 1.05, y_{t-2} = 1.10$, even though the slope of both triangular planes in y_{t-2} is zero. Finally, rising up from the lower left-hand plane, to connect with the higher right-hand plane, is a steeply sloped logistic curve that runs along the diagonal of the forecasting surface

from the lower front corner to the top back corner in the region where $y_{t-1} \approx y_{t-2}$. This steep connecting section is the area in which the nonlinear ANN term is "active." The importance of each of the three forecasting regions in Figure 13 is discussed below.

We now explain the meaning of the short vertical pillars rising up from the floor of Figure 13. To do this, recall that, as explained in Section 2, the functional relationship plotted in Figure 13 is used to forecast a stream of y s from September 1929 out into the (almost) infinite future, using only market data from August 1929 and earlier, and that the present value of this forecasted y stream is used to estimate September 1929's fundamental price in Equation (3). As described in Appendix 2, this is accomplished by first forecasting y_t given actual data on y_{t-1} , y_{t-2} and a randomly obtained u innovation. We next use the same forecasting function to forecast y_{t+1} using as inputs into our model the just-forecasted y_t , the actual value of y_{t-1} , and another randomly drawn innovation. In similar fashion, forecasted values of y_t and y_{t+1} are next used to forecast y_{t+2} , and so forth, until a forecast for $y_{t+10,000}$ is obtained. We then repeat this entire forecasting procedure 9,999 times—using each time a new sequence of η disturbances, but the same forecasting function, in standard Monte Carlo fashion—so that we obtain 10,000 different forecasts for the sequence $y_t, y_{t+1}, y_{t+2}, \dots, y_{t+10,000}$. As explained earlier, the multiplier $E_{t-1} \left\{ \sum_{k=0}^{\infty} \left(\prod_{i=0}^k y_{t+i} \right) \right\}$ is then calculated as the average value of the sum, $\sum_{k=0}^{10,000} \left(\prod_{i=0}^k y_{t+i} \right)$, from our 10,000 economies.

Notice from Equation (3) that forecasted y_t appears in every product on the right side of Equation (3), while forecasted y_{t+1} appears in all but one product, y_{t+2} in all but two products, etc. Early values of forecasted y_{t+i} therefore have a much greater effect on the entire product sum in Equation (3) than do later values of forecasted y_{t+i} . For this reason, the behavior of the first few y_{t+i} forecasted largely determines the value of $E_{t-1}\{\cdot\}$ in Equation (3). In Figure 13 we therefore plot, as the 60,000 short vertical pillars rising up from the figure's floor, the first 6 months of Monte Carlo evolutions for y_{t+i} (measured on the vertical axis) as a function of y_{t+i-1}, y_{t+i-2} (i.e., the two most recently forecasted values of y in the Monte Carlo sequence) from our 10,000 simulated economies produced in September 1929. The five tall pillars that extend from the floor of Figure 13 all the way up through the forecasting surface represent, from top to bottom, the maximum, 99.5th percentile, median, 0.5th percentile, and minimum values of y_{t+i} forecasted in September 1929. Notice that most of the pillars occur along the diagonal of the forecasting surface in the area occupied by the ANN logistic curve. This provides visual confirmation that the ANN terms are indeed important for forecasting y .

With all the elements of Figure 13 accounted for, we can finally explain why the ANN terms are important for forecasting discounted dividend growth. To do this, recall that the upper triangular plane in Figure 13 is situated higher up the y_t axis than is the lower triangular plane; that is, for a given value of y_{t-1} , the forecasted value of y_t is higher for low values of y_{t-2} than it is for a high values of y_{t-2} , all else constant. Also note that the logistic curve rises from the lower plane, where $y_{t-1} < y_{t-2}$, up to the upper plane, where $y_{t-1} > y_{t-2}$, as y_{t-2} falls below y_{t-1} . Thus, unlike the two triangular regions of the forecasting surface, which have zero slope in the y_{t-2} direction, the region along the diagonal of Figure 13 is negatively sloped in y_{t-2} . The diagonal region is also more positively sloped in y_{t-1} than are the two triangular planes (which have the same slope as each other), since starting from a point on the lower plane where $y_{t-1} < y_{t-2}$, and increasing y_{t-1} until $y_{t-1} > y_{t-2}$, results in an upward move to the higher plane.

The fact that Figure 13's forecasting surface is more steeply sloped in both directions when $y_{t-1} \approx y_{t-2}$ than it is when either $y_{t-1} \ll y_{t-2}$ or $y_{t-1} \gg y_{t-2}$, as in the triangular regions, is important for explaining the ANN's success in forecasting dividends. To see why, notice that most of the y pillars protruding up from the floor of Figure 13 pierce the forecasting surface along the surface's diagonal; that is, in the area occupied by the logistic ANN curve. Thus, in the majority of Monte Carlo iterations, $y_{t+i-1} \approx y_{t+i-2}$ so that y_{t+i} is forecasted from (y_{t+i-1}, y_{t+i-2}) with a large positive weight on y_{t+i-1} (approximately 1.7) and a smaller negative weight on y_{t+i-2} (approximately -0.9). However, when y_{t+i-1} and y_{t+i-2} are far apart, we move to Figure 13's flatter triangular planes in which the slope in y_{t+i-1} is only 0.85 and the slope in y_{t+i-2} is zero. The nonlinear ANN term therefore allows us to treat different pairs of y_{t+i-1}, y_{t+i-2} differently.

Unlike our flexible ANN model, a standard linear ARAR or ARMA model is forced to give the same lag weights to past y values no matter what their behavior or history. Thus, sudden and potentially explosive outlier y s are viewed in the same functional manner as subtle changes in y 's behavior. A standard ARMA model therefore compromises between its desire to capture small subtle changes and its aversion to outlier overreaction by selecting a more gently sloped forecasting surface that, by necessity, mutes its response to subtle y innovations. Conversely, our ANN model is able to allow for more dramatic responses to subtle y movements in the $y_{t+i-1} \approx y_{t+i-2}$ region, where almost all observations lie, while simultaneously avoiding overreaction to outliers by having a different functional relationship in the outlier region. The ANN's increased modeling flexibility in the $y_{t+i-1} \approx y_{t+i-2}$ region leads to more persistence in y_{t+i} 's simulated

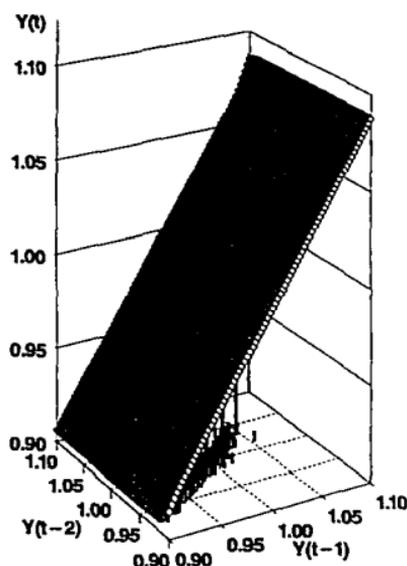
evolution in the vast majority of cases. For reasons explained in the previous section on ARCH effects, this increased persistence in simulated y series leads to higher fundamental prices for the ANN-included Figure 7 than for the no-ANN Figure 10.

3.3 The behavior of y and the timing of the peak

We have thus far seen that, to obtain fundamentals that mimic the general rise and fall in market prices, the approach employed must be sufficiently flexible to have the potential to fit the data (Section 3.1) and the model used must be well specified (Section 3.2). In this section we discuss factors that are especially important for our ability to capture the timing of the September 1929 peak in market prices. In particular, we discuss why it is that our bond yield- r fundamental price turns down in October 1929 at exactly the same time when the stock market actually crashed. More importantly, we discuss why it is that our fundamental price peaks in 1929 instead of in other years, such as 1927, when the dividend growth rate was even higher.

We begin by investigating the October 1929 decline in Figure 7's fundamental price. To do this, we plot in Figure 14 the forecast mapping surface, and the 60,000 pillars that represent the first 6 months of forecasts from each of our 10,000 simulation replications, that are produced by the bond yield- r model from October 1929. The construction of Figure 14 is analogous to that of Figure 13 described above. Notice that Figure 14's forecasting surface appears virtually identical to Figure 13's surface. The similarity of these two surfaces reveals that, although we update our estimates of model parameters every month as the investors' information set is updated, the parameter values for the October 1929 forecasting function are virtually identical to those from September 1929. This finding suggests that the ability of Figure 7's bond yield- r model to fit the turning point in market prices is not due to a sudden and substantial shift in parameter values between September and October, but is instead due to an important change in the nature of y 's (lagged) behavior between September and October 1929. This observation is confirmed by noticing the substantial downward shift (i.e., a shift toward the origin of the graph) from Figure 13 to Figure 14 in the tall pillars which represent the forecasts of y_{t+i} . This downward shift in the distribution of the simulated economies' forecasted discounted dividend growth rates, y_{t+i} , produces an $E_{t-1}\{\cdot\}$ multiplier that is smaller in October 1929 than in September 1929 (see Figure 9) and thus yields a fundamental price in October 1929 that is lower than the September 1929 fundamental (see Figure 7).

We now proceed to the discussion of why our fundamental price series peaks in 1929, instead of in some other year, by plotting in Figure 15 the bond yield- r 's y_t series; that is, the time series of realized



Forecasting Surface: Diamonds

Data Points: Pillars

Figure 14

Forecasting function: October 1929

This figure plots, as the surface of diamonds, the functional relationship between (y_{t-1}, y_{t-2}) and y_t implied by our model in Table 2, with the discount rate r_t equal to the interest rate on low-risk commercial paper plus a constant equity premium equal to the average excess return on stocks over low-risk bonds from 1871 to 1988. Parameter values are from October 1929, the first month after the stock market crash. The 60,000 short vertical pillars rising up from the figure's floor represent the first 6 months of Monte Carlo evolutions for y_{t+i} (measured on the vertical axis) as a function of y_{t+i-1}, y_{t+i-2} (i.e., the two most recently forecasted values of y in the Monte Carlo sequence) from our 10,000 simulated economies produced in October 1929. The five tall pillars that extend from the figure's floor all the way up through the forecasting surface represent, from top to bottom, the maximum, 99.5th percentile, median, 0.5th percentile, and minimum values of y_{t+i} forecasted in October 1929.

discounted dividend growth rates (i.e., the data) on which our model is estimated and forecast simulations are based. Notice that, while y_t reaches a local maximum in mid-1929, the 1929 peak in y is not a global maximum. Indeed, the y_t series peak in 1929 is somewhat lower than the 1927 peak in y , though the 1929 fundamental price in Figure 7 is considerably higher than the 1927 fundamental.

The 1929 fundamental price peak is higher than the 1927 price peak for two reasons. First, as seen in Figure 3, the level of dividends was higher in 1929 than it was in 1927. Thus, even if the $E_{t-1}\{\cdot\}$ dividend multiplier in 1929 was equal to the 1927 multiplier, as in the basic Gordon model, the fundamental price (i.e., dividends times the multiplier) would be higher in 1929. Of course, as seen directly in Figure 9, our model's $E_{t-1}\{\cdot\}$ dividend multiplier is considerably

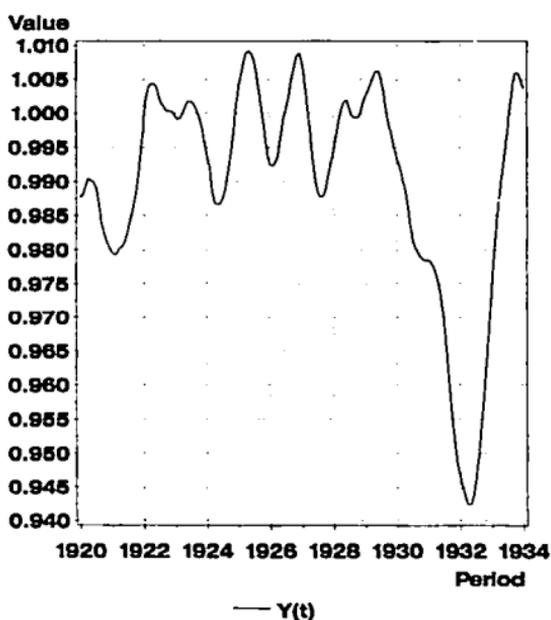


Figure 15
Discounted dividend growth rate (bond yield-r)

This figure plots the discounted dividend growth rate, $y_t = (1 + g_t)/(1 + r_t)$, where g_t is the dividend growth rate and r_t is the interest rate on low-risk commercial paper plus a constant equity premium equal to the average excess return on stocks over low-risk bonds from 1871 to 1988. See Appendix 1 for data sources.

higher in 1929 than it was in 1927, which is the second reason the 1929 price is higher; the increased $E_{t-1}\{\cdot\}$ multiplier magnifies 1929's increased dividend levels to produce the 1929 price peak.

There are two reasons why we obtain a 1929 dividend multiplier $E_{t-1}\{\cdot\}$ that is larger than the multiplier from 1927, or any previous year. First is the declining ARCH variance effect, discussed in Section 3.2.2, which causes more weight to be placed on the autocorrelated component of our model and less weight to be placed on the u_t innovations, as we enter the late 1920s. The second, previously undiscussed, reason is that the autocorrelated component of our model becomes more autocorrelated during the late 1920s. This feature is seen directly in Table 3, which reports parameter estimates and associated standard errors for every second December from 1919 to 1934. Note the downward shift in the regression constant α and the shift upward in the AR1 parameter β , both from Equation (6), which occurs in mid-1927. This α - β switching implies that the data becomes more persistent during the late 1920s than it had previously been. Thus, at the very moment that Section 3.3.2's ARCH effects begin placing more weight on the persistent component of y , y becomes more persistent.

Table 3
Parameter estimates for equations (6) through (12): bond yield- r (standard errors in parentheses)

Equation number→	(6)	(6)	(7)	(10)	(10)	(12)	(12)
Parameter name→	α	β_1	δ_1	ρ_1	ρ_2	λ	ξ_1
Date ↓							
1919:12	0.438* (0.094)	0.558* (0.095)	-2.619* (0.648)	1.708* (0.081)	-0.775* (0.075)	0.372* (0.071)	0.202 (0.141)
1921:12	0.477* (0.086)	0.519* (0.087)	-2.765* (0.658)	1.742* (0.066)	-0.800* (0.063)	0.338* (0.070)	0.194 (0.153)
1923:12	0.503* (0.082)	0.493* (0.083)	-2.860* (0.653)	1.760* (0.060)	-0.817* (0.057)	0.326* (0.069)	0.198 (0.155)
1925:12	0.516* (0.082)	0.479* (0.083)	-2.930* (0.675)	1.765* (0.058)	-0.820* (0.056)	0.318* (0.071)	0.234 (0.162)
1927:12	0.158 (0.237)	0.839* (0.238)	-3.403* (0.930)	1.347* (0.290)	-0.463* (0.131)	0.290* (0.070)	0.250 (0.171)
1929:12	0.148 (0.197)	0.849* (0.198)	-3.746* (0.864)	1.319* (0.238)	-0.440* (0.099)	0.267* (0.070)	0.257 (0.183)
1931:12	0.135 (0.171)	0.861* (0.172)	-3.968* (0.931)	1.343* (0.219)	-0.442* (0.107)	0.265* (0.068)	0.224 (0.182)
1933:12	0.120 (0.173)	0.876* (0.175)	-4.615* (0.143)	1.325* (0.212)	-0.416* (0.096)	0.243* (0.067)	0.242 (0.191)

This table contains estimated parameter values for the ARAR-ARCH-ANN forecasting model in Table 2. The discount rate employed in the formation of y_t is the interest rate on riskfree debt plus a constant equity premium, as explained in Appendix 1. Although the model is reestimated every month, only parameter estimates from December in odd numbered years are reported due to space constraints. Throughout the table, * denotes statistical significance at 5 percent.

The multiplicative combination of increased dividend growth persistence, and increased weight being placed on this persistence, leads to superpersistence in the evolution of forecast simulation y_{t+i} s during the late 1920s; a superpersistence we call the "Roaring '20s Effect." Indeed, the emergence of this superpersistence in y_t 's evolution suggests that late 1920s investors could have rationally believed that any changes in dividend behavior would have a more lasting effect on future dividends than had previously been the case. In other words, when the news about dividend growth (i.e., y_t) was good, late 1920s investors would be much more optimistic than their early 1920s counterparts. Similarly, when dividend news was bad, the behavior of the dividend data suggests that late 1920s investors would fundamentally be more pessimistic than their early 1920s counterparts. Changes in

dividend growth that had produced only small price responses in the early 1920s would therefore fundamentally lead to much larger price responses in the late 1920s and early 1930s.

Along with all of the specification and modeling factors discussed earlier in this section, recognition of the Roaring '20s Effect goes a long way toward explaining the market price boom and crash. To again quote White (1990:72), price-dividend plots such as Figure 3 "reveal the remarkable change that overtook the stock market [during the late 1920s]. From 1922 to 1927 dividends and prices moved together, but while dividends continued to grow rather smoothly in 1928 and 1929, stock prices soared far above them." While, to the naked eye, this 1927 switch in behavior may appear to imply a sudden bubble of overoptimism, our more thorough analysis offers a fundamental data-driven explanation. In particular, our results suggest that the relationship between prices and dividends changed during the late 1920s because the time-series behavior of discounted dividend growth changed. Indeed, our findings imply that, while an investor living in 1920 might view a sudden change in discounted dividend growth as a predominantly temporary shock, an investor living in the late 1920s would be more easily persuaded to view a sudden increase in y as the harbinger of persistently increasing prosperity. This finding is consistent with claims made by economic historians that investors living in the roaring '20s believed they had entered a new age of peace and prosperity and therefore viewed economic events with unfettered optimism. Our results reveal that, while ex post dividend realizations were less than expected, the 1920s optimism may have been rational *at the time*, given the observed behavior of discounted dividend growth.

3.4 The effect of various discount rate assumptions

In the preceding sections we have focused primarily on results obtained with the discount rate—that is, r_t in $y_t = (1 + g_t)/(1 + r_t)$ —defined as the real return on high-grade short-term debt plus a constant equity premium, as in Figure 7. We now investigate why our model, with r_t defined as a constant discount rate, as in Figure 5, peaks a little later than the bond yield- r fundamental in Figure 7, and why defining $(1 + r_t)$ as a consumption-based discount rate, as in Figure 6, yields a fundamental price that peaks later still.

The time-series behavior of the constant- r and consumption- r y_t series are in general very similar to the behavior of the bond yield- r y_t series plotted in Figure 15, as can be inferred from the statistics in Table 1. Indeed, the main reason we did not plot all three series in Figure 15 is that the scale of the figure is not fine enough for the naked eye to adequately distinguish one series from the other. How-

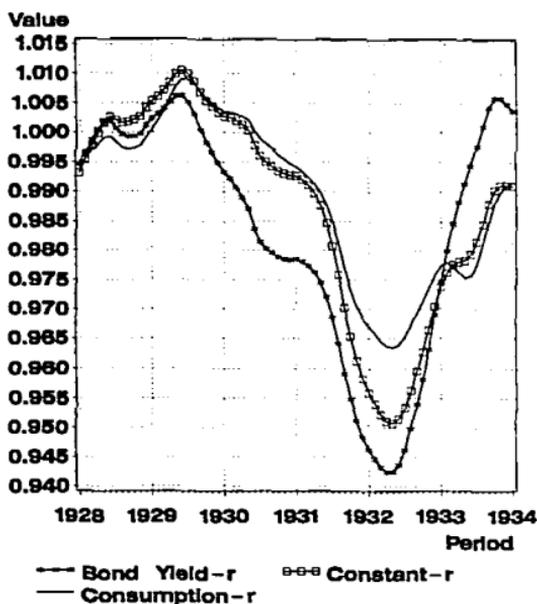


Figure 16
Discounted dividend growth rates: $Y(t)$

This figure plots the discounted dividend growth rate, $y_t = (1 + g_t)/(1 + r_t)$, under various assumptions for r_t . For the line joining squares, the discount rate r_t is a constant equal to the average return on stocks from 1871 to 1988. For the line joining stars, the discount rate r_t is the interest rate on low-risk commercial paper plus a constant equity premium equal to the average excess return on stocks over low-risk bonds from 1871 to 1988. For the plain line, the gross discount rate $(1 + r_t)$ is derived from a standard power utility function in consumption with the coefficient of relative risk aversion $\alpha = 1.5$ and the monthly subjective time discount rate $\beta = 0.9953$ (i.e., 0.945 annually). See Appendix 1 for data sources.

ever, what small differences do exist between the three series are obviously sufficient to produce perceptible differences in fundamentals, especially around the time of the crash in 1929. To focus on this important turning point, we plot in Figure 16 all three y_t series from 1928 to 1932.

Before embarking on a detailed discussion of Figure 16, note that any variation in $y_t = (1 + g_t)/(1 + r_t)$ for the constant- r series must be due to variation in $(1 + g_t)$, since $(1 + r_t)$ is constant by assumption. The observation that all three y_t series in Figure 16 are very similar reveals that the vast majority of variation in the bond yield- r and consumption- r y_t series is also due to movements in dividend growth, as opposed to movements in the discount rate, and thus that broad movements in the fundamental prices these series produce (see Figures 5 through 7) are mostly due to dividend behavior. Conversely, one can argue that finer issues, such as timing of the peak, are influenced in important ways by the discounting method employed. This

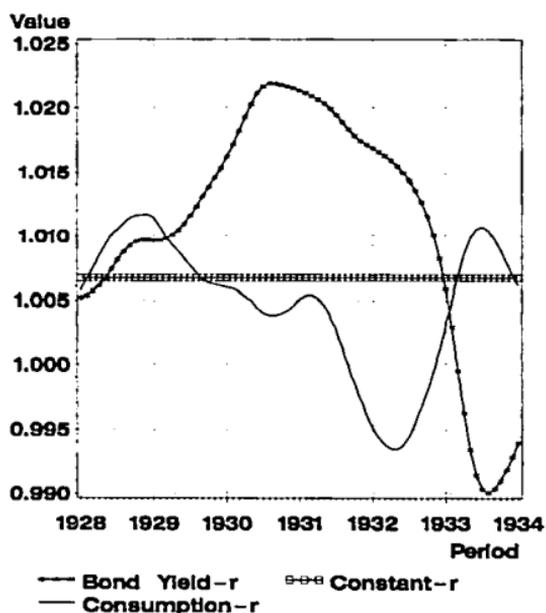


Figure 17
Gross discount rates

This figure plots the discount rate $(1 + r_t)$ under various assumptions. For the horizontal line joining squares, the discount rate r_t is a constant equal to the average return on stocks from 1871 to 1988. For the line joining stars, the discount rate r_t is the interest rate on low-risk commercial paper plus a constant equity premium equal to the average excess return on stocks over low-risk bonds from 1871 to 1988. For the plain line, the gross discount rate $(1 + r_t)$ is derived from a standard power utility function in consumption with the coefficient of relative risk aversion $\alpha = 1.5$ and the monthly subjective time discount rate $\beta = 0.9953$ (i.e., 0.945 annually). See Appendix 1 for data sources.

argument is examined in Figure 17, which plots the $(1 + r_t)$ denominator from y_t from the three discounting assumptions.

Consider first the constant- r y_t series in Figure 16, which is the highest line (joining squares) at the 1929 peak. Notice that the constant- r y_t series, and thus the discounted dividend growth rate, reaches its maximum in July 1929 and then falls somewhat rapidly from late 1929 until the middle of 1932. Of course, as revealed in Figure 3's plot of the dividend level (multiplied by 200), the level of dividends continued to rise—although obviously at a declining rate—until early 1930. However, after September 1929 this sustained increase in dividend levels was more than offset by the dramatically falling dividend multiplier whose value is determined by the declining y_t in Figure 16. The net result of these two counterbalancing forces—rising dividends and falling dividend multiplier—is therefore a constant- r fundamental in Figure 5 that peaks only 3 months late.

Now consider, in comparison, the consumption- r series represented

by the plain line in Figures 16 and 17. Looking at Figure 17, we see that the denominator of consumption- r y_t , and thus the consumption ratio C_t/C_{t-1} , had already started to decline by the beginning of 1929 and kept falling until the middle of 1932. This decline in y_t 's denominator retarded y_t 's descent from 1929 to 1932, as is evidenced by the fact that the plain consumption- r line in Figure 16 falls more slowly than any other y_t line over this period. This slower fall in the consumption- r y_t causes our autoregressive forecasting model to produce a slowly falling forecasted y series, which in turn produces a consumption- r dividend multiplier that falls more slowly than the constant- r multiplier. Since dividends were still rising through early 1930, the consumption- r y_t 's slow decline explains why the consumption-based fundamental in Figure 6 peaks 20 points too high and 7 months late.

Finally, consider the bond yield- r case represented by the line joining stars in Figures 16 and 17. This fundamental fits market prices more closely than our other models in part because, as seen in Figure 16, the bond yield- r y_t falls farther and faster than any other y series. The bond yield y_t also rises farther and faster in 1933, which explains why the bond yield- r fundamental in Figure 7 mimics the market recovery more closely than the other fundamentals in 1933. The reason for the bond yield- r fundamental price's superior performance, relative to the constant- r and consumption- r fundamental series, is evident in Figure 17. From Figure 17 we see that the bond yield- r 's denominator for y_t rises through the crash of 1929 to peak in late 1930, thus magnifying the concomitant fall in y_t 's numerator ($1 + g_t$), as seen in Figure 16. Similarly, just as dividend growth starts to rise in late 1932, the bond yield- r is falling, again magnifying the short-term fluctuation in dividend growth. While from Figure 16 (and from a comparison of Figures 5 and 7) it is obvious that differences between the constant- r and bond yield- r discounting assumptions do not drive our general ability to reject the bubbles hypothesis, these small differences do influence second-order effects, such as the exact timing of the peak.

3.5 Summary

In the preceding discussion we have documented four key factors in the production of our results. First, we have seen in Section 3.1 that to obtain our results we must use an unrestricted form of the present value relationship between expected future discounted dividends and current asset prices, as shown in Equation (3). Indeed, we have seen that Gordon's familiar constant dividend growth restriction is sufficiently at odds with the data that models built on this assumption cannot reproduce the full boom and bust in 1920s stock prices. Second, we have seen in Section 3.2 that, to accurately capture key features of

the data necessary to reject bubbles, one must employ a time-series model for discounted dividend growth, y_t , that does not fail standard specification tests, including tests for uncaptured variance effects and nonlinearity. In particular, we have shown that failure to account for ARCH and ANN effects in the data yields a fundamental price that can fail to reject bubbles. Third, we have seen in Section 3.3 that, to capture subtle changes in the way the market interprets movements in the data, model parameters—including variance effects—must be continuously updated to include the most current information available to investors. Specifically, we have shown the importance of accounting for what might be called the “Roaring ‘20s Effect” as dividend growth becomes more persistent and accurately forecastable with the progression of time. Finally, we have seen in Section 3.4 that our ability to fit the exact timing of the peak depends in part on the discount rate assumption we employ. We favor the bond yield- r assumption because, as explained at the beginning of Section 2, this particular specification allows us to work with ratios of nominal discount and growth rates directly and thus removes a possible source of measurement error in the price index used to form inflation. However, even with a constant- r or consumption- r convention, it seems difficult to claim that there was a bubble in the 1929 stock market.

4. Concluding Remarks

In this article we have introduced a new procedure for estimating fundamental stock prices as the present value of expected future cash flows. Our procedure differs from those currently employed in two key respects. First, instead of focusing on dividend levels or dividend growth alone, we chose the discounted dividend growth rate series y_t as our object of interest. Second, instead of assuming that dividends grow at a constant rate from each forecasting date into the future, we use more flexible time-series techniques and Monte Carlo simulation to forecast future dividend paths conditional on information available to investors at the time stock prices were actually being set in the market. In particular, the model specifications and parameter estimates we employ in our forecasting exercise are determined using only in-sample data and not by mining or snooping the out-of-sample data to find models that make sense *ex post*. However, we do undertake a rather thorough search and conduct several diagnostic tests to ensure that our forecasting models contain the many elements required to be well specified in-sample.

To examine the potential of our new procedure, we have studied the relationship between market prices and estimated fundamentals during the Great Stock Market Crash of 1929. The traditional quantita-

tive evidence in support of the "conventional wisdom" that the 1920s stock market contained a bubble relies on the observations that dividends grew more slowly than prices during the late 1920s and that the expost warranted price, based on realizations of post-1929 dividends, does not share the spectacular rise and fall in actual prices. However, we have argued that restrictions on the behavior of dividends and dividend growth that are necessary to obtain the expost or Gordon prices as "fundamental" expectations of future dividends, conditional only on information available to investors before the boom and crash, are not consistent with the pre-boom and crash dividend data. Thus, tests of the present value model based on expost warranted and Gordon prices do not provide reliable evidence of a bubble in 1920s stock prices. They simply reveal misspecifications in the traditional fundamentals generating procedures employed.

Conversely, we have shown that our more general ARAR-ARCH-ANN models of the discounted dividend growth process do capture key features of the discounted dividend growth data, including a time-varying mean and variance and important nonlinear effects. Using our models, and only data on dividends, bond yields, and consumption available to investors at the time prices were actually being determined in the market, we have produced fundamental prices that match the magnitude and timing of the boom and crash in 1929 stock prices. Statistical tests confirm that our fundamental prices also share important time-series properties with actual prices and reject the hypothesis that market prices contain a bubble. We therefore conclude that, given the information available to investors living in the early 1920s, dividends may well have been expected to increase by enough to warrant the observed rise in market prices and that, as new information arrived in the late 1920s, expectations of future dividends were revised downward resulting in the observed crash in prices. Thus, although we can never be sure exactly what market participants expected in the way of future dividends, it does appear that there is at least one reasonable fundamental explanation for the boom and crash in 1920s stock prices.

Our investigation of this single historical episode, the crash of 1929, has an important implication for the usefulness of present value relationships in studying the behavior of asset prices in general. As stated in the introduction, the boom and crash in 1920s stock prices is often used as an extreme example of the failure of the traditional bubble-free present value model. However, our analysis rejects the "conventional wisdom" and instead reveals that there is indeed a fundamental explanation for at least this one event. Our results therefore weaken the case for bubbles in general and suggest that a very cautious and thorough investigation of the data should be undertaken

before labeling even the most spectacular gyration in market prices as an event at odds with the standard present value model.

Finally, recall from our discussion in Section 2.1 that to forecast future discounted dividend growth we have used in this study only information on past discounted dividend growth. In particular, we have not used past market prices to forecast future dividends for fear that, if there truly was a bubble in the market price, then we could have inadvertently imputed a bubble into our dividend fundamentals via past market price information and thus erroneously failed to reject the bubble hypothesis. Given that we have demonstrated an absence of bubbles with a forecasting information set that excludes past market prices, however, it now seems safer to include a much wider variety of information—such as past market prices, earnings, macro factors, etc.—in the forecasting information sets employed in future applications of our procedure. Indeed, the use of larger information sets in the application of our procedure to a more extensive analysis of present value relationships is the subject of ongoing research.

Appendix 1: Data

All monthly data are collected from 1899:01 to 1934:12. The years 1899 and 1934 are used for lead/lag purposes so that all monthly data used in estimations are from 1900:01 to 1933:12. The following series are employed.

Stock Prices: S&P500 stock price index monthly from Cowles (1939) series P1.

Producer Prices/Inflation: Monthly producer price index from Macaulay (1938). As one would expect, this series contains a strong seasonal pattern at monthly and quarterly frequencies and also exhibits occasional spikes caused by sudden and often temporary shocks to the index's component factors. To account for these features, we first deseasonalize the data with the standard X11 procedure and then smooth out the remaining transitory spikes using a standard spline function. Inflation rates from the resulting monthly prices q are calculated as $\pi_t = (q_t - q_{t-1})/q_{t-1}$.

Interest Rates: For calculations based on a constant discount rate, we use an annual constant real rate of 8.3 percent, which equals the average annual real return on stocks from 1871 to 1988 [data from Shiller (1989)]. The monthly constant real discount rate is therefore $r = 1.083^{1/12} - 1 = 0.0067$. When discounting cash flows at a variable riskless rate plus constant risk premium, we use as our riskless return the rate on 4 to 6 month prime commercial paper from Macaulay (1938), filtered using the procedure described above for producer prices. The monthly nominal riskless interest rate, R , is calculated from

the annual rate, R_a , as $R = (1 + R_a)^{1/12} - 1$ and the monthly real interest rate is calculated as $r = (R - \pi)/(1 + \pi)$. To this we add a constant real monthly risk premium of 0.0040. This is the monthly equivalent of a 4.9 percent real annual premium, which is the average real annual premium of stocks over prime commercial paper between 1871 and 1988 calculated with annual data from Shiller (1989). Calculations are performed using Shiller's numbers on the S&P 500 (series 1), S&P 500 dividends (series 2), riskless interest rate (series 4), and producer prices (series 5). The monthly risk premium ρ is then calculated from the annual numbers as $\rho = (1 + \rho_a)^{1/12} - 1$. Our annual premium of 4.9 percent differs from the 6 percent reported by Mehra and Prescott (1985) because of the different time periods studied (they used 1899 to 1978) and because they used a consumption deflator instead of producer prices. We use producer prices because the consumption deflator is not available on a monthly basis.

Dividends: Monthly dividends on the S&P 500 index are obtained jointly from series C^1 , stock prices including cash dividends, and P^1 , stock prices, from Cowles (1939) as $D_t = (P_t^1 C_{t+1}^1 / C_t^1) - P_{t+1}^1$. Since, like producer prices and interest rates, dividends are seasonal and spiked, we also filter dividends with the procedure outlined above for prices and interest rates. To make sure that our filtering process has not destroyed important properties of the data, we calculated annual dividends from both the raw and filtered monthly series, where the annual dividend is the compounded sum of the dividends from January to December of each year. (This is equivalent to a single December lump payment; i.e. $D_{annual} = \sum_{i=1}^{12} (\prod_{j=1}^i [1 + r_j]^{-1j-1}) D_i$.) The average difference between the two annualized dividend series is a miniscule \$0.0006, on an average \$3.98 annual dividend (or about $\frac{1}{6500}$ of the dividend amount), and displays no strong time-series pattern. We found even less of a difference between the raw and filtered interest rate and producer price series. This leads us to conclude that our filter has not destroyed important properties of the data as they pertain to our calculations of fundamental stock prices.

Consumption: Consumption is not available monthly for our time period, but is available on a quarterly basis from 1919 onward in Balke and Gordon (1986). However, industrial production [from Miron and Romer (1990)] and dividends (see above) are both available monthly over our entire 1899 to 1934 time period and, according to models such as Lucas (1978), are both at least theoretical proxies for consumption. We therefore formed a consumption proxy in the following manner. First, we deseasonalized quarterly real consumption, in the manner described above for prices, and formed the quarterly ratio C_t / C_{t-3} , which we then cube-rooted and placed in each month for the applicable quarter from 1919 to 1934. We then regressed this

consumption ratio series on the monthly real dividend ratio and real production ratio for 1919 to 1934 to obtain a monthly functional relationship between consumption, dividends, and production. We then used this estimated functional relationship, along with dividend and production data, to obtain a monthly consumption ratio proxy for each month from 1899 to 1934.

Appendix 2: Technical Details

Monte Carlo Simulation

First, all parameter values in Equations (6) through (13) are estimated on the 240 monthly observations from 1900:01 to 1919:12. Each of the 240 resulting in-sample u_{t-i} residuals ($i = 1, \dots, 240$) from Equation (10) are then divided by the corresponding date $t-i$ estimate of $\sqrt{b_{t-i}}$ ($i = 1, \dots, 240$) from Equation (12) to produce a time series of 240 i.i.d. η_{t-i} ($i = 1, \dots, 240$) innovations (e.g., u_{t-1} is divided by $\sqrt{b_{t-1}}$ to produce η_{t-1} , u_{t-2} is divided by $\sqrt{b_{t-2}}$ to produce η_{t-2} , and so forth). These 240 η_{t-i} standardized residuals form the pool of i.i.d. innovations used in subsequent Monte Carlo work.

Second, we use in-sample information from 1900:01 to 1919:12 to obtain an out-of-sample forecast for b_t in Equation (12). Third, we randomly draw from our i.i.d. η pool a value for the η_t Monte Carlo innovation and then multiply this random η_t by our forecasted $\sqrt{b_t}$ to obtain a bootstrapped u_t . We then insert this bootstrapped u_t into Equation (10) along with past values of ϵ_{t-i} to produce a simulated ϵ_t . Values for z_{t-i} are then obtained from Equation (8) using past y_{t-i} , inserted into Equation (7) to obtain $\Psi(\cdot)$, which is then inserted into Equation (6) along with the simulated ϵ_t and past values of y_{t-i} to produce a simulated value for y_t . This simulated y_t becomes the first element inside the parentheses in Equation (3).

Fourth, we update by one period steps 2 and 3 above (but not step 1) to create a simulated value for y_{t+1} , ϵ_{t+1} , u_{t+1} , and b_{t+1} using the just-simulated values for y_t , ϵ_t , u_t , and b_t , a randomly drawn η_t , and actual data for variables dated $t-1$ and earlier (note that we still use *actual* data from before period t only). The simulated y_{t+1} is then multiplied by the simulated y_t to form the second element in the parentheses in Equation (3). We continue to roll out simulated values for y_{t+2} , y_{t+3} , y_{t+4} , and so forth, using steps two and three and actual data from before period t , until the product of the forecasted y s falls below 0.00001. (Since forecasted—and actual— y s usually assume values that are on average just under one, this convergence criterion typically involves simulating y out for from 5,000 to 10,000 periods; that is, roughly 400 to 800 years into the future.) Fifth, we calculate

the sum of the progressive product of the simulated y s to obtain a simulated estimate of the parenthesized term on the right-hand side of Equation (3). This ends our first loop through the simulation process.

Step 6 involves repeating the entire simulation/roll-out procedure described in steps 2 through 5 for a new series of randomly drawn η s. This second trip through the simulation loop yields a second estimate of the parenthesized term on the right-hand side of Equation (3). This loop is repeated until 10,000 simulated estimates of the parenthesized term on the right-hand side of Equation (3) are obtained for the data and parameter estimates from 1919:12. In step seven we then take the average of our 10,000 simulated sums of products of y s to yield $E_{1919:12}\{\cdot\}$ in Equation (3). This average is finally multiplied by the dividend from 1919:12 to yield our fundamental price for 1920:1, $P_{1920:1}^F$.

We then update our information set in rolling window fashion so that the 240 observations ending in 1920:1, and the entire simulation procedure from steps 1 through 7, are used to obtain a fundamental price for 1920:2. Subsequent fundamental prices are produced in the same updated forecast simulation fashion until the entire fundamental price series for 1920 to 1933 is completed.

Model Selection

Our exact model selection routine is as follows. Note from the grid specified in Equation (13) that the system [Equations (6) through (13)] contains $3^7 = 2187$ possible specifications. On the presample data for 1900 to 1919, we estimated for the constant- r , consumption- r , and bond yield- r y processes candidate specifications. Specification tests, such as heteroskedasticity-robust tests for residual autocorrelation and LM tests for uncaptured ARCH, were then conducted and all specifications that did not pass these tests were discarded. The remaining models were then ranked according to their ability to minimize the Schwarz model selection criterion. Of all well-specified models, the five specifications that produced the best Schwarz criterion values were selected for further investigation. For each of the constant- r , consumption- r , and bond yield- r y s, we used each of our five Schwarz-best models to generate simulated y sequences using steps 1 through 4 of the Monte Carlo procedure described above and *only data from before 1919*. The models that produced simulated y s that had properties closest to the actual sequence of y s from 1900 to 1919 (e.g., similar means, variances, etc.) were finally selected as our ultimate model specifications for use in the forecasting of y out-of-sample for 1920 to 1933.

References

- Ackert, L., and B. Smith, 1993. "Stock Price Volatility, Ordinary Dividends, and Other Cash Flows to Shareholders," *Journal of Finance*, 48, 1147-1160.
- Balke, N., and R. Gordon, 1986. "Appendix B: Historical Data," in R. Gordon (ed.), *The American Business Cycle: Conformity and Change*, University of Chicago Press, Chicago.
- Barsky, R., and B. DeLong, 1993. "Why Does the Stock Market Fluctuate?" *Quarterly Journal of Economics*, 108, 291-311.
- Camerer, C., 1989. "Bubbles and Fads in Asset Prices," *Journal of Economic Surveys*, 3, 3-38.
- Campbell, J. Y., and R. J. Shiller, 1987. "Cointegration Tests of Present Value Models," *Journal of Political Economy*, 95, 1062-1088.
- Cochrane, J., 1992. "Explaining the Variance of Price-Dividend Ratios," *Review of Financial Studies*, 5, 243-280.
- Cowles, A., 1939. *Common Stock Indexes*, Principia Press, Bloomington, Indiana.
- Dickey, D., and W. Fuller, 1979. "Distribution of the Estimates for Autoregressive Time Series with a Unit Root," *Journal of the American Statistical Association*, 74, 427-31.
- Flood, R., and P. Garber, 1994. *Speculative Bubbles, Speculative Attacks and Policy Switching*, MIT Press, Cambridge, Mass.
- Flood, R., and R. Hodrick, 1990. "On Testing For Speculative Bubbles," *Journal of Economic Perspectives*, 5, 124-150.
- Garber, P., 1990. "Famous First Bubbles," *Journal of Economic Perspectives*, 4, 35-54.
- Gordon, M., 1962. *The Investment, Financing and Valuation of the Corporation*, Irwin, Homewood, Ill.
- Grossman, S., and R. Shiller, 1981. "The Determinants of the Variability of Stock Prices," *American Economic Review*, 81, 222-227.
- Hornik, K., M. Stinchcombe, and H. White, 1989. "Multi-layer Feedforward Networks are Universal Approximators," *Neural Network*, 2, 15-38.
- Hornik, K., M. Stinchcombe, and H. White, 1990. "Universal Approximation of an Unknown Mapping and its Derivatives Using Multilayer Feedforward Networks," *Neural Network*, 3, 551-560.
- Kindleberger, C. P., 1978. *Manias, Panics and Crashes: A History of Financial Crisis*, Basic Books, New York.
- Kleidon, A., 1986. "Variance Bound Tests and Stock Price Valuation Models," *Journal of Political Economy*, 94, 953-1001.
- Kuan, C., and H. White, 1994. "Artificial Neural Networks: An Econometric Perspective," *Econometric Reviews*, 13, 1-91.
- Lee, T., H. White, and C. W. J. Granger, 1993. "Testing for Neglected Nonlinearity in Time Series Models," *Journal of Econometrics*, 56, 269-290.
- Lucas, R. E., Jr., 1978. "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1429-1445.
- Macaulay, F., 1938. *The Movement of Interest Rates, Bond Yields and Stock Prices in the United States Since 1858*, National Bureau of Economic Research, New York.

Mankiw, G., D. Romer, and M. Shapiro, 1985, "An Unbiased Reexamination of Stock Market Volatility," *Journal of Finance*, 40, 677-687.

Mankiw, G., D. Romer, and M. Shapiro, 1991, "Stock Market Forecastability and Volatility: A Statistical Appraisal," *Review of Economic Studies*, 58, 455-477.

Mehra, R., and E. Prescott, 1985, "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 16, 115-134.

Miron, J., and K. Romer, 1990, "A New Monthly Index of Industrial Production 1884-1940," *Journal of Economic History*, 50, 321-337.

Phillips, P. C. B., and P. Perron, 1988, "Testing for a Unit Root in Time Series Regression," *Biometrika*, 75, 335-46.

Shiller, R., 1981, "Do Stock Prices Move too Much to be Justified by Subsequent Changes in Dividends?" *American Economic Review*, 71, 421-436.

Shiller, R., 1989, *Market Volatility*, MIT Press, Cambridge, Mass.

Stinchcombe, M., and H. White, 1994, "Using Feedforward Networks to Distinguish Multivariate Populations," forthcoming in *Proceedings of the International Joint Conference on Neural Networks*.

West, K., 1987, "A Specification Test for Speculative Bubbles," *Quarterly Journal of Economics*, 102, 553-580.

West, K., 1988, "Bubbles, Fads and Stock Price Volatility Tests: A Partial Evaluation," *Journal of Finance*, 43, 639-656.

White, E., 1990, "The Stock Market Boom and Crash of 1929 Revisited," *Journal of Economic Perspectives*, 4, 67-84.