NBER WORKING PAPER SERIES

INTRINSIC BUBBLES: THE CASE OF STOCK PRICES

Kenneth A. Froot

Maurice Obstfeld

Working Paper No. 3091

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 1989

The authors are grateful to John Campbell, Steve Durlauf, Jeff Frankel, Greg Mankiw, Jeff Miron, Andy Rose, Julio Rotemberg, Jeremy Stein, and especially Jim Stock for helpful comments. We also received useful suggestions from participants-especially the discussant, Ken West-at the 1989 FMME/Summer Institute conference. Bob Barsky and Brad DeLong helped us obtain data. We are also grateful to the John M. Olin, Alfred P. Sloan, and National Science Foundations for generous financial support, and to the IMF Research Department for its hospitality while this draft was completed. This paper is part of NBER's research program in Financial Markets and Monetary Economics. Any opinions expressed are those of the authors not those of the National Bureau of Economic Research.

NBER Working Paper #3091
September 1989

INTRINSIC BUBBLES: THE CASE OF STOCK PRICES

ABSTRACT

Several puzzling aspects of the behavior of United States stock prices can be explained by the presence of a specific type of rational bubble that depends exclusively on dividends. We call such bubbles "intrinsic" bubbles because they derive all of their variability from exogenous economic fundamentals, and none from extraneous factors. Unlike the most popular examples of rational bubbles, intrinsic bubbles provide an empirically plausible account of deviations from present-value pricing. Their explanatory potential comes partly from their ability to generate persistent deviations that appear relatively stable over long periods.

Kenneth A. Froot
Harvard Business School
Dillon 33
Soldiers Field
Boston, MA 02162

Maurice Obstfeld
Department of Economics
Harvard University
Cambridge, MA 02138

1. Introduction

After a decade of research, financial economists remain unable to account for the temporal volatility of stock prices. The initial rejections by LeRoy and Porter (1981) and Shiller (1981) of simple present-value models based on constant discount rates and rational expectations have been weakened, but not reversed, by subsequent work. Departures from present-value prices still appear large and persistent.[1]

At one time rational bubbles were viewed as a promising alternative hypothesis. Interest in this alternative has waned, however, because econometric tests have not produced strong positive evidence that rational bubbles can explain asset prices. That is, no one has produced a specific bubble parameterization which is both simple yet capable of explaining the data.

In this paper we propose and test empirically a new rational-bubble specification with both these properties. Our formulation is simple because it introduces no extraneous sources of variability. Instead, the bubbles we examine are driven exclusively – albeit

[1] Campbell and Shiller (1987), Flavin (1983), Froot (1988), Kleidon (1986), Mankiw, Romer, and Shapiro (1985), Marsh and Merton (1986), and West (1987, 1988a) address econometric shortcomings of the original studies. Attempts to extend the simple present-value model to allow for time-varying discount rates have added little; see Campbell and Shiller (1988a, 1988b), Flood, Hodrick, and Kaplan (1986), and Shiller (1981). Pindyck (1984) suggests that low-frequency price fluctuations may be a result of time-varying risk premia driven by changing stock-price volatility. However, Poterba and Summers (1986) argue that volatility is not sufficiently persistent to explain a large portion of low-frequency price movements.

nonlinearly – by the exogenous fundamental determinants of asset prices. For this reason we refer to these bubbles as "intrinsic."[2] One striking property of intrinsic bubbles is that, for any given level of exogenous fundamentals, asset prices remain constant over time: intrinsic bubbles are deterministic functions of fundamental alone. Thus, this class of bubbles predicts that stable and highly persistent fundamentals lead to stable and persistent over- or undervaluations.

Intrinsic bubbles also appear capable of explaining long-term movements in stock prices. It turns out that the component of prices not explained by the present-value model is highly correlated with dividends, as an intrinsic bubble would predict. These bubbles therefore capture the apparent overreaction of prices to dividend changes. For example, they appear to explain the bull market of the 1960s, a period of high and rising real dividends, as well as the market decline of the early 1970s. We use our estimated model to separate out the present-value and bubble components of stock prices, and find that the former implies a realized return on stocks of about 9.1 percent – very close to the 9.0 percent average for this century.

Of course, other alternative non-bubble hypotheses could conceivably explain our results. It is well known that any bubble path is observationally equivalent to a present-value path where the process generating fundamentals may change in the future.[3] Our results could thus be interpreted in principle as evidence of such prospective changes. Indeed, in an exchange-rate model with stochastic regime changes, we have derived present-value pricing formulas similar to the bubble formulas derived below.[4] In this paper, however, we offer no particular regime-switch model to explain the apparently nonlinear relationship

[2] The excessive variability of an intrinsic bubble solution comes entirely from its functional form, not from the introduction of extraneous state variables. In models with stationary sunspot equilibria, asset prices generally can be expressed as functions of fundamentals alone. However, some of these fundamentals, real interest rates in particular, are endogenous and at least one ultimate source of their variability could be an extraneous state variable. An intrinsic bubble solution for an asset price is a reduced-form expression that depends only on the exogenous factors affecting the economy, not on extraneous noise. In other words, every intrinsic bubble solution is a "minimal-state-variable" solution in the sense of McCallum (1983).

[3] Flood and Garber (1980), Hamilton and Whiteman (1985), and Flood and Hodrick (1986) discuss this observational equivalence. Cecchetti, Lam, and Mark (1989) study the empirical properties of particular nonlinear-fundamentals-forcing processes.

[4] See Froot and Obstfeld (1989).

between prices and dividends.

A second alternative hypothesis is that stationary fads or noise lie behind the departures from present-value prices.[5] Both fads and intrinsic bubbles predict that these departures will be highly persistent. But an important theoretical distinction between the two is that the former entail short-term speculative profit opportunities, whereas bubbles alone do not. In our empirical tests, the intrinsic bubble formulation allows us to identify separately these two sources of deviation from present-value pricing. While the predictability of short-term returns may ultimately be useful in explaining certain features of the data, our results suggest that this predictability is not the main explanation for the present-value model's failure.

The paper is structured as follows. Section 2 shows how intrinsic bubbles arise in a standard present-value model. We compare in section 3 some properties of intrinsic and extraneous bubbles. Section 4 then turns to the data. We examine the univariate and bivariate times-series properties of U.S. stock prices and dividends, and argue that an intrinsic bubble is broadly consistent with the results. In the second part of section 4, we estimate our model directly and test it against several alternatives. Section 5 concludes.

2. Intrinsic bubbles in a present-value model

Stochastic linear rational expectations models can have a multiplicity of solutions that depend on exogenous fundamentals but do not depend on extraneous variables such as time.[6] In this section we describe how such rational bubbles can arise as nonlinear solutions to a linear asset-pricing model. Although our choice of a specific model is guided by the empirical application we have in mind, solutions similar to those derived below arise in a broader class of models.

The model is based on the standard arbitrage condition linking the time series of real

[5] For examples of models with fads or noise, see Black (1986), Campbell and Kyle (1988), DeLong, Shleifer, Summers and Waldman (1988), Shiller (1984), and Summers (1986).

[6] Included in the category of extraneous variables would be irrelevant fundamentals, such as lagged fundamentals that play no economic role apart from their effect on expectations.

stock prices to the time series of exogenous dividend payments. Let P_t be the real price of a share at the beginning of period t, D_t real dividends per share paid out over period t, and r the constant, instantaneous real rate of interest. The arbitrage condition we focus on is

$$P_t = e^{-r} E_t(D_t + P_{t+1}), \tag{1}$$

where $E_t(.)$ is an expectation conditional on information known at the start of period t.[7]

The present-value solution for P_t, denoted by P_t^{pv}, is:

$$P_t^{pv} = \sum_{s=t}^{\infty} e^{-r(s-t+1)} E_t(D_s). \tag{2}$$

Equation (2) is a particular solution to the stochastic difference equation (1). It equates a stock's price to the present discounted value of expected future dividend payments. We will always assume that the present value (2) exists, that is, that the continuously compounded growth rate of expected dividends is less than r.

The present-value formula is the solution to (1) usually singled out by the relevant economic theory as a unique equilibrium price. It can be derived by applying the transversality condition,

$$\lim_{s \to \infty} e^{-rs} E_t(P_s) = 0, \tag{3}$$

and then observing that successive forward substitutions into (1) converge to (2).

Equation (1) has solutions other than (2). By construction, these alternative price paths satisfy the requirement of period-by-period efficiency, but they do not satisfy (3). Let $\{B_t\}_{t=0}^{\infty}$ be any sequence of random variables such that

$$B_t = e^{-r} E_t(B_{t+1}). \tag{4}$$

Then $P_t = P_t^{pv} + B_t$ is a solution to (1), which can be thought of as the sum of the present-value solution and a rational bubble. Clearly, property (4) implies that P_t violates the transversality condition (3).

[7] In our empirical implementation of the model below, we allow for errors in the arbitrage equation.

Rational bubbles are typically viewed as being driven by variables extraneous to the valuation problem. However, some bubbles may depend only on the exogenous fundamental determinants of asset value. We call such bubbles "intrinsic" because their dynamics are inherited entirely from those of the fundamentals. An intrinsic bubble is constructed by finding a nonlinear function of fundamentals that satisfies (4). In the above stock-price model with only one stochastic fundamental factor – the dividend process – intrinsic rational bubbles depend on dividends alone.

To see what an intrinsic stock-price bubble might look like, suppose that log dividends are generated by the geometric martingale,

$$d_{t+1} = \mu + d_t + \xi_{t+1}, \tag{5}$$

where μ is the trend growth in dividends, d_t is the log of dividends at time t, and ξ_{t+1} is a normal random variable with conditional mean zero and variance σ^2.[8] Using (5), and assuming that period-t dividends are known when P_t is set, we see that the present-value stock price in (2) is directly proportional to dividends,

$$P_t^{pv} = \kappa D_t, \tag{6}$$

where $\kappa = (e^r - e^{\mu+\sigma^2/2})^{-1}$. Equation (6) is essentially a stochastic version of the Gordon (1962) model of stock prices, which predicts that $P_t^{pv} = (e^r - e^{\mu})^{-1} D_t$ under certainty. The assumption that the sum in (2) converges implies that $r > \mu + \sigma^2/2$.

Now define the function $B(D_t)$ as

$$B(D_t) = cD_t^{\lambda}, \tag{7}$$

where λ is the positive root of the quadratic equation

$$\lambda^2\sigma^2/2 + \lambda\mu - r = 0, \tag{8}$$

[8] Kleidon (1986) uses this specification in his empirical study.

and c is an arbitrary constant. It is easy to verify that (7) satisfies (4):

$$e^{-r}E_t(B(D_{t+1})) = e^{-r}E_t(cD_t^{\lambda}e^{\lambda(\mu+\xi_{t+1})}) \tag{9}$$

$$= e^{-r}(cD_t^{\lambda}e^{\lambda\mu+\lambda^2\sigma^2/2}) = e^{-r}(cD_t^{\lambda}e^{r}) = B(D_t).$$

By summing the present-value price and the bubble in (7), we get our basic stock-price equation:

$$P(D_t) = P_t^{pv} + B(D_t) = P_t^{pv} + cD_t^{\lambda}. \tag{10}$$

Even though (10) contains a bubble (for $c \neq 0$), and thus violates (3), it is driven exclusively by fundamentals: $P(D_t)$ is a function of dividends only, and does not depend on time or any other extraneous variable. $B(D_t)$ is therefore an example of an intrinsic bubble.[9]

The inequality $r > \mu + \sigma^2/2$ can be used to show that λ must always exceed 1. It is this explosive nonlinearity that permits $B(D_t)$ to grow in expectation at rate r. We will assume from now on that $c > 0$, so that stock prices cannot be negative. Negative stock prices would violate free disposability.[10]

It might seem paradoxical that movements in a bubbly asset price can be accounted for completely by movements in fundamentals. Economists are accustomed to an almost instinctive decomposition of asset prices into two components, one dependent on market fundamentals, and a second reflecting self-fulfilling beliefs and driven, at least in part, by extraneous factors. In the context of linear models, for example, McCallum (1983) argues that bubble solutions can be avoided by restricting attention to "minimal-state-variable" solutions that depend only on fundamentals. The possibility of intrinsic bubbles reveals that McCallum's approach does

[9] Sargent (1987, pp. 348-349) characterizes a rational bubble as a function $\bar{B}(t, X_t) = e^{rt}X_t$ of time and a variable X_t that obeys $E_t(X_{t+1}) = X_t$. His definition does not imply, however, that bubbles have to contain deterministic time components. To write the bubble $B(D_t)$ defined by (7) in Sargent's form, simply let $X_t = e^{-rt}cD_t^{\lambda}$.

[10] Let λ' be the negative root of equation (8). Then the general solution to (1) (within the class of functions $P = P(D_t)$) is:

$$P(D_t) = P_t^{pv} + c_1D_t^{\lambda} + c_2D_t^{\lambda'}.$$

We have imposed $c_2 = 0$ in (10) on the grounds that the stock price P_t should go to zero (not to infinity) as dividends D_t go to zero. The argument in the text shows that *any* variable Y_t whose logarithm follows a martingale with drift μ and variance σ^2 leads to a bubble solution to (1), $P(D_t, Y_t) = P_t^{pv} + B(Y_t)$. Thus, a formula like (7) can be used to construct extraneous as well as intrinsic bubbles.

not rule out multiple solutions unless some additional requirement – linearity of the price function, for example – is imposed.[11]

Like all rational bubbles, intrinsic bubbles rely on self-fulfilling expectations. Instead of being driven by extraneous variables, however, these expectations are driven by the nonlinear form of the price solution itself. Figure 1 shows the family of solutions (10) for a particular choice of $c > 0$. The straight line $P^{pv}P^{pv}$ indicates the present-value solution (6); this solution implies that $E_t(P_{t+1}/P_t) = e^{\mu+\sigma^2/2} < e^r$. A point like 1 on the bubble path satisfies the arbitrage condition (1) because of Jensen's inequality. At point 1, the next innovation in the log of dividends is distributed symmetrically around zero, but the market's belief that the relevant price function has the shape shown means that the expected rise in the stock price, and hence the stock price itself, is higher at point 1 than at the corresponding point 2 on $P^{pv}P^{pv}$.[12]

3. Intrinsic versus extraneous bubbles: A partial comparison

Why should one think that intrinsic bubbles might succeed in characterizing asset prices when other bubble formulations have failed? In this section we argue that intrinsic bubbles have several intuitively appealing properties which are absent in the bubble parameterizations used previously in empirical studies.

To begin, we need to know why bubble explanations of stock prices have fared so poorly.[13] A first reason might be a belief that prices simply do not diverge from their present-value levels. There certainly are theoretical arguments for holding this view, but it has proven difficult empirically to reconcile observed price behavior with a wide range of present-value models. The theoretical conditions required to rule out rational bubbles

---

[11] Another counterexample comes from models in which calendar time itself is a fundamental. Then deterministic time-driven bubbles of the Flood and Garber (1980) sort satisfy the minimal-state-variable criterion.

[12] It is easy to check that various theorems used to identify *unique* solutions of the form $P(D_t)$ to equations like (1) do not apply under this section's assumptions. For example, (10) is not within any of the classes of solutions considered by Lucas (1978), Saracoglu and Sargent (1978), Gourieroux, Laffont, and Monfort (1982), or Whiteman (1983). The problem is not that the process (5) is nonstationary. Multiple solutions analogous to (10) exist when (5) is a mean-reverting Ornstein-Uhlenbeck process; see Froot and Obstfeld (1989). Rather, the problem is that standard solutions impose additional restrictions, such as linearity of the solution or the assumption that all state variables are restricted to assume values in compact sets. These assumptions rule out solutions such as (10).

[13] Flood and Hodrick (1989) present a detailed survey of the empirical literature on bubbles.

are relatively demanding; these conditions assume substantial, perhaps unrealistic, long-horizon foresight on the part of economic agents. Short-horizon excess-profit opportunities, on the other hand, are plausibly quite small.

A second reason for the poor empirical track record of bubbles is that the specific parameterizations that have been tested have also failed. These parameterizations generally assume that bubbles depend explicitly on time.[14] As a result, they predict upward runs in stock prices conditional on dividends. There is little evidence, however, either for price runs or for price-dividend ratios trending deterministically upward through time. These features of the data suggest that time-driven bubble formulations are too restrictive to improve our understanding of asset prices.

Some general specification tests have been employed in the hope of detecting bubbles, without taking a stand on a specific bubble form. Even though these tests may have low power, they nevertheless reject the no-bubble null frequently. However, they cannot reveal the precise source of rejection, so they yield no hard evidence that bubbles really are the culprit.[15] The tendency to ascribe these rejections to sources other than bubbles has been strengthened, perhaps excessively, both by the theoretical arguments against bubbles and the failure of the specific parameterizations mentioned above. However, consideration of stochastic bubbles that look quite different from the typical time-driven examples may throw a different light on the specification-test results.

How then do intrinsic bubbles look, and why might they do a better job of explaining prices? First, intrinsic bubbles capture well the idea that stock prices overreact to news about dividends, as argued by Shiller (1984), among others. Equation (10) implies that $\frac{dP_t}{dD_t} = \kappa + \lambda c D_t^{\lambda-1} > \kappa$, so that prices move more when dividends change than the present-value formula (6) would predict.

[14] See Flood and Garber (1980) and Blanchard and Watson (1982) for specific examples.

[15] The general specification test for bubbles used by West (1987) and Casella (1988) can alternatively be interpreted as a test of model specification, the purpose for which it was originally proposed by Cumby, Huizinga, and Obstfeld (1983). A second type of specification test for bubbles compares the time-series properties of prices and dividends, which should differ if condition (1) holds but stock prices contain a rational bubble. See Hamilton and Whiteman (1985) and Diba and Grossman (1988a).

Intrinsic bubbles may also help explain the time-series behavior of prices. Even though prices are predicted to grow at the rate of interest, specific realizations may fluctuate within some limited range for rather long periods. A given dividend realization corresponds to a unique stock price regardless of the date on which the dividend is announced. Because dividends are persistent, deviations from present-value prices may also be highly persistent. An implication of this property is that, even with a very long data series, the fundamentally explosive nature of an intrinsic bubble might be impossible to detect econometrically.

To illustrate these points, we present some simulations comparing the intrinsic bubble in (10) with a particular alternative bubble specification. Each simulation experiment involves three solutions to the difference equation (1). The first of these is the present-value price P_t^{pv} given by (6); the second is a nonlinear intrinsic bubble of the form (10), denoted by $\hat{P}_t$; and the third is a bubble that depends on time as well as on dividends,

$$\tilde{P}_t = P_t^{pv} + bD_t e^{(r-\mu-\sigma^2/2)t}. \tag{11}$$

The precise formulation in (11) is chosen for two reasons: First, it makes the bubble a function of dividends, and thus allows stock prices to overreact to dividend news, just as the bubble (10) does. Second, (11) follows the majority of parametric bubble tests in adopting a specification in which the extraneous variable t affects prices.

Dividends are assumed to follow (5), and in each experiment successive innovations ξ_t are drawn independently from a normal distribution. P_t^{pv} is calculated using estimates of r, μ, and σ^2 extracted from U.S. stock-price and dividend data, and the values of the parameters κ, c, and b are those estimated below in section 4. The simulations are run over 200 years. However, it is important to note that there is little importance to these specific choices of parameters and sample size: the qualitative patterns displayed in the following figures are quite general.

Figure 2 shows a first run in which the simulated intrinsic bubble, $\hat{P}_t$, does not produce noticeable explosive behavior within the simulation sample. The percentage overvaluation of stocks is not very different at the end of the sample (the year 2100) than it is around 1970 or 2015. In contrast, the time-driven bubble $\tilde{P}_t$ explodes decisively starting in mid-sample.

The behavior of the time-driven bubble is similar in Figure 3, but the underlying dividend realization makes the explosive expected growth of the intrinsic bubble more apparent. Figures 2 and 3 highlight the sharply different paths for intrinsic bubbles that different paths of fundamentals can produce.

Diba and Grossman (1988b) have argued on theoretical grounds that stochastic rational bubbles cannot "pop" and subsequently start up again. This feature, they assert, makes rational bubbles empirically implausible. Figure 4, however, shows an intrinsic bubble realization that falls over time to a level quite close to fundamentals. Indeed, if dividends follow a process like (5) but without drift, the logarithm of dividends reaches any given lower bound in finite time with probability one; and we can therefore be sure that the bubble term in (10) gets arbitrarily close to zero in finite time. For practical purposes, this is the same as periodically popping and restarting with probability one. Intrinsic bubbles can get very close to present-value prices, and then diverge.

Notice that all three simulations share the common feature that the intrinsic bubble path lies above the time bubble in the early part of the sample, but below it by the sample's end. This pattern in the early part of the sample is merely a result of initial conditions, and is therefore purely arbitrary.[16] By contrast, the feature that the time bubble eventually exceeds the intrinsic bubble is more general. It is easy to show as the sample size t grows, the probability that $\hat{P}_t > \tilde{P}_t$ goes to zero, for any set of initial conditions.[17] Although the intrinsic bubble $\hat{P}_t$ ultimately exceeds the time-driven bubble $\tilde{P}_t$ very rarely in large samples, when it does, it does so by an amount large enough to equalize the two bubbles' expected growth rates.

[16] It turns out that if model (11) is to have any hope of fitting the data, the estimate of b must be very close to zero, implying that $\tilde{P}_t$ is very close to P_t^{pv} for the first part of the sample. See Section 4.2 and Figure 6 below.

[17] *Proof:* Define $\psi = r - \mu - \sigma^2/2$ and assume, without loss of generality, that the bubbles are equal at $t = 0$: $bD_0 = cD_0^\lambda$. Then $Prob[\tilde{P}_t < \hat{P}_t] = Prob[bD_t e^{\psi t} < cD_t^\lambda] = Prob[\psi t < (\lambda - 1)(\mu t + \sum_{s=1}^t \xi_s)] = Prob[r - \lambda\mu - \sigma^2/2 < (\sum_{s=1}^t \xi_s)/t]$. Equation (8) implies, however, that $r - \lambda\mu - \sigma^2/2 = \sigma^2(\lambda^2 - 1)/2 > 0$ (recall that $\lambda > 1$). Since $plim(\sum_{s=1}^t \xi_s)/t = 0$, the proof is complete.

This latter property is important empirically. It implies that it would be unusual to draw a long dividend series which yields an intrinsic bubble that appears as explosive as a comparable time-driven bubble. Even though intrinsic and time-driven bubbles are expected to grow at the same rate on average, a long intrinsic-bubble sample path is very likely to appear less explosive than the path a time-driven bubble such as (11) generates.

4. Application to the U.S. stock market

In this section we turn to U.S. stock market data to examine the empirical performance of our model. In doing so, we generalize slightly the model in section 1 to allow for errors in the initial arbitrage equation. Thus, time-t prices are now given by:

$$P_t = e^{-r} E_t(D_t + P_{t+1}) + e^{-r} u_t, \tag{1'}$$

where u_t is an error term, assumed to be independent of dividends at all leads and lags and to have unconditional mean zero.[18]

Equation (1′) allows us to express (10) as the statistical model,

$$P_t = c_0 D_t + c D_t^{\lambda} + \epsilon_t, \tag{12}$$

in which $c_0 = \kappa = (e^r - e^{\mu+\sigma^2/2})^{-1}$ and ϵ_t is the present value of the errors in (1′), $\epsilon_t \equiv \sum_{s=t}^{\infty} e^{-r(s-t+1)} E_t(u_s)$. The error u_t is a predictable single-period excess return, and ϵ_t is its infinite-horizon counterpart. These excess returns could be interpreted, for example, as the result of time-varying effective income tax rates, provided that those rates are conditionally independent of D_t.[19] One could also think of u_t as partly reflecting a fad – a shock to the demand for stocks which is unrelated to efficient forecasts of future

[18] This distributional assumption is unnecessarily strong. Our tests below will produce consistent parameter estimates provided only that $E_t(\xi_{t+j}|u_t) = 0 \ \forall j \geq 0$. The evidence in the first appendix below supports this assumption. The stronger assumption that $E_t(\xi_{t+j}|u_t) = 0 \ \forall j$ is needed for consistent inferences.

[19] More formally, suppose that dividends are taxed at the marginal rate θ_t at time t, and that the rate follows the process

$$\theta_{t+1} = \rho\theta_t + \omega_{t+1},$$

with $0 \leq \rho < 1$ and with θ_t and D_t independently distributed at all leads and lags. It is then easy to show that the present discounted value of future dividend *receipts* satisfies $P_t^{pv} = \kappa D_t + \epsilon_t$, where $\epsilon_t = -\rho(e^r - \rho e^{\mu+\sigma^2/2})^{-1}\theta_t D_t$, and $E(\epsilon_t|D_s) = 0 \ \forall s$. Furthermore, this step leads to specification (13) below.

dividends. Thus, our empirical specification allows us to identify separately bubble and fad components of stock prices.[20]

Estimation of (12) is complicated by collinearity among the regressors. To mitigate this problem, we divide (12) by D_t to express the price-dividend ratio as a nonlinear function of dividends:

$$\frac{P_t}{D_t} = c_0 + cD_t^{\lambda-1} + \eta_t, \tag{13}$$

where the new error term, $\eta_t = \epsilon_t/D_t$, satisfies $E(\eta_t|D_s) = 0$, $\forall s$. The null hypothesis of no bubble implies that $c_0 = \kappa$ and $c = 0$; whereas the bubble alternative in (10) predicts that $c_0 = \kappa$ and $c \neq 0$.

In the estimation below we use the Standard and Poor's stock-price and dividend indexes from the *Securities Price Index Record*, as extended backwards in time by Cowles *et al.* (1939). Following Barsky and DeLong (1989), we examine the period 1900-88, using nominal stock prices recorded in January of each year and deflated by the January PPI. Dividends are annual averages of nominal data for the calendar year, deflated by the year-average PPI.[21] Of course, we would like to have data on beginning-of-period-t dividends to match the beginning-of-period-t price, P_t. Because these are not available, we use the average of period-t dividends as our measure of D_t.[22][23]

4.1. The price-dividend relation

In deriving (13), we assumed that the log dividend process follows a martingale with

[20] Campbell and Shiller (1987, 1988a, 1988b), for example, rule out rational bubbles from the start, and therefore attribute deviations from present-value pricing entirely to a stationary fad component.

[21] Although the price and dividend series have been extended back to 1871, we chose to begin our sample at 1900 for two reasons. First, the composition of the market portfolio becomes increasingly restrictive as one goes back in time. By the 1870s the portfolio is comprised of just 11 railroad stocks. Second, whereas January values for the PPI are available after 1900, only annual averages exist prior to 1900. Because many other authors (e.g., Campbell and Shiller 1987) have used the longer series, we also ran all of the statistical tests below on the 1871-1986 sample. The results were qualitatively unaffected.

[22] A potential problem with this choice is that D_t may not be completely known at the beginning of period t. Nevertheless, we see two reasons why D_t is likely to be a better measure of the dividend information contained in beginning-of-period-t price, P_t, than is the average period-t-1 dividend, D_{t-1}. First, P_t is not recorded on January 1, but is itself an average over the period-t month of January. Second, to mitigate the effects of any time lapse between the determination and actual distribution of dividends, it is better to use average period-t dividends than those from period $t-1$. In any case, unless otherwise mentioned the results below are not importantly different when average period-$t-1$ dividends are used to proxy for beginning-of-period-t dividends.

[23] In applying our specification to the aggregate stock market, a natural question is how an intrinsic bubble dependent on aggregate dividends could arise. One possibility is that each firm's share price equals the present value of its own dividends, plus an intrinsic bubble on *aggregate* dividends. (This would require that an individual firm's dividends do not Granger cause aggregate dividends.) Such a formulation would remove the incentive for managers to influence the market price of their firms' shares by altering the timing of dividend payments.

trend. While this particular stochastic process is chosen for simplicity, we wish to be sure that it is at least a reasonable approximation to actual dividend behavior. In the first appendix below, we describe several univariate and bivariate tests of the log dividend specification in (5). We find little evidence against the martingale hypothesis: log dividend changes are essentially unpredictable when conditioning on the lags of log dividends and/or log price-dividend ratios.[24] The data estimate the parameters in (5) as $\mu = 0.011$ and $\sigma = 0.122$.

A general implication of (13) is that stock prices may appear to overreact to changes in dividends. Also, (13) predicts that price-dividend ratios are nonstationary and positively correlated with dividends. This subsection presents a brief empirical examination of these basic implications of intrinsic bubbles.

First, what does the present-value model predict for the sensitivity of prices to changes in dividends? From (6) a one dollar change in dividends should raise prices by κ dollars. Using the fact that the the sample-average gross real return on stocks is $e^r = 1.090$ per annum, we have that $\kappa = (e^r - e^{\mu+\sigma^2/2})^{-1} = (1.090 - e^{.011+.122^2/2})^{-1} = 14.0$. In general if P_t and D_t are cointegrated of order $(1,1)$, then under the present-value model the cointegrating coefficient should be approximately κ. Equation (6) also implies that the elasticity of prices with respect to dividends is 1. If the log stock price p_t and d_t are cointegrated, it is also with a coefficient of 1.

The first line of Table 1 presents estimates of κ, obtained by regressing prices on dividends. The coefficient is estimated to be 36.7 – much larger than the value of 14.0 predicted by our present-value model.[25] If P_t and D_t are cointegrated then the OLS estimate of the cointegrating factor, while consistent, is biased in small samples. In order to bound the cointegrating coefficient, we run the reverse regression – projecting D_t on P_t – in the second line of Table 1. This produces an even larger estimate of κ, $1/.0233 = 42.9$.

[24] Some of this evidence may be controversial. We have placed our discussion in the first appendix because the controversial aspects are tangential to our main argument.

[25] Similar estimates of the cointegrating factor are obtained by Campbell and Shiller (1987), Diba and Grossman (1988a), and West (1987), among others.

These values seem too large to be consistent with the present-value model. Even the lower of the two would imply that the required rate of return on stocks less the expected growth rate of dividends is an implausibly low $1/36.7 = 2.7$ percent per annum. (The actual value over our sample period is 7.1 percent.) The third and fourth lines of Table 1 perform analogous regressions in logs instead of levels. Here the cointegrating coefficient predicted by the present-value model is 1, but the estimates are again much higher – between 1.59 and $1/.5563 = 1.80$. These estimates suggest that simple present-value models cannot explain why price-dividend ratios are so high given historical stock returns, or, equivalently, why returns have been so high given price-dividend ratios.

To test whether these estimates are statistically incompatible with the present-value model, we examine various measures of the price-dividend ratio for nonstationarity. Table 2 reports Phillips-Perron unit root tests for the theoretically warranted "spread", $P_t - 14D_t$, as well as the price-dividend ratio in levels, P_t/D_t, and in logs, $p_t - d_t$. Results of tests with and without time trends are reported. Under the present-value model, we should reject nonstationarity in each of these regressions. Yet in five of six cases we cannot reject the unit-root hypothesis. Of course, the power of these tests may be low, but the evidence for stationarity seems too ambiguous to justify ruling bubbles out by assumption.[26]

In sum, the evidence presented in this section has three important implications for our argument. First, prices are too sensitive to current dividends to be consistent with a simple present-value model. The implication, of course, is that the portion of stock prices unexplained by such a model must be highly correlated with dividends.[27] Second

[26] Some of our results may be sensitive to the timing of dividends. Diba and Grossman (1988a), for example, use lagged dividends and deflate by the WPI. They find that the log price-dividend ratio, $p_t - d_{t-1}$, is stationary. Using lagged dividends, but deflating by the PPI, Campbell and Shiller (1988a) also reject nonstationarity. Tests using lagged dividends, however, may reject too frequently under the assumption that $p_t - d_t$ actually contains a unit root. Campbell and Shiller (1987) find results similar to those reported above for the spread, $P_t - \kappa D_t$, using data from 1871-1986. All of these authors acknowledge that the evidence is not clear cut, but their maintained assumption that price-dividend ratios are stationary is critical to interpreting their findings.

[27] We are certainly not the first to notice this fact, which is essentially a robust restatement of Shiller's (1981) volatility findings. More specifically, West's (1987) general specification test and Campbell and Kyle's (1988) noise trading model exploit the excess sensitivity of prices to dividend changes. Durlauf and Hall (1988) find noise in prices that is more highly correlated with prices themselves than with dividends. Their definition of noise, however, is not the difference between prices and a multiple of current dividends, but the difference between prices and an ex post measure of the present value of future dividends.

this overreaction cannot be explained by other variables which are incorporated into stock prices and which help forecast future dividends. If, for example, when dividends are high investors tend to get other reliable information that dividends will grow more quickly than previously expected, then this information is likely to be incorporated in stock prices, which therefore should Granger-cause dividends. The results in the first appendix suggest, however, that this is the not case. Finally, a specification such as (13) has at least the potential to explain these failures of the present-value model.

4.2. A direct test for intrinsic bubbles

To see if this potential is at all realized, we turn in Table 3 to estimates of (13) and several related expressions. Before interpreting the estimates, however, some discussion of econometric issues is in order.

The regressor in (13) presents difficulties because it is explosive. Two additional assumptions are necessary for valid statistical inferences. If the t-statistic of $c = 0$ is to have the usual distribution we require that: i) The residuals, η_t, are distributed normally – but not necessarily identically or independently – with unconditional mean zero; and ii) the dividend innovations, ξ_t, are distributed independently of the residuals η_t at all leads and lags. The second appendix provides a proof that the standard t-statistic does indeed approximate a normal distribution under these assumptions, despite the presence of the exploding regressor, $D^{\lambda-1}$.

The other aspect of estimation that requires discussion is estimation of the standard error of the residual η_t. Because theory gives us no guide to η_t's serial correlation, the usual standard errors may be incorrect. We try to account for this possibility in two ways. First, we estimate (13) by OLS, but correct the residuals using Newey and West's (1987) covariance-matrix estimator for serial correlation of unknown form. This estimator also allows for conditional heteroskedasticity.[28] Second, since the residuals appear to be

[28] It is sensible to think of the residual in (12), ϵ_t, as growing at a rate similar to that of dividends (see footnote 19). In such a case, we would not expect η_t to exhibit much conditional heteroskedasticity. Indeed, in our estimates the heteroskedasticity-corrected standard errors were similar to the uncorrected standard errors.

well described by a first-order autoregressive process, we compute maximum likelihood estimates of the parameters under the assumption that the residuals are AR(1).

Finally, there is the issue of how to estimate the exponent, λ, and the present-value multiplier, κ. In some of the regressions below we do not estimate λ concurrently with the other parameters. Instead, we use the point estimates from the log dividend process obtained earlier, together with the mean return on stocks over the period to compute $\lambda = 2.74$.[29] In other regressions we do estimate all parameters simultaneously, without imposing additional restrictions. The restriction that $c_0 = \kappa = 14.0$ is not imposed on the constant term in (13), even though it holds under both the null and alternative hypotheses. Instead, we use the unrestricted estimate of c_0 as a kind of sensibility check on our model.

The first two lines of Table 3 report estimates of (13) using OLS and maximum likelihood, respectively. These two regressions constrain λ to equal 2.74. In both cases, $\hat{c}$ is statistically very significant. The estimates are comparable in magnitude and significance for the two estimation methods.[30] In the third and fourth lines we estimate all of the parameters of the model simultaneously. The point estimates of c_0 are similar to those above, although λ is estimated to be larger and c correspondingly lower.[31] The larger standard error for c is expected here because the estimates of c and λ are highly collinear.[32] Rather than using a t-test to judge the importance of the nonlinear term, it is therefore more appropriate to compute an F-test of the no-bubble hypothesis, $c = 0, \lambda = \hat{\lambda}$, where $\hat{\lambda}$ is the unrestricted estimate of λ reported in the third and fourth lines, respectively. This

[29] We tried a variety of parameter estimates for r, μ, and σ^2. These do have a minor effect on the exponent but are unimportant for the general regression results reported below.

[30] We also tried estimating an extended form of (13),

$$\frac{P_t}{D_t} = c_0 + c_1 D_t^{\lambda - 1} + c_2 D_t^{\lambda' - 1} + \eta_t,$$

where λ' is the negative root from equation (8). Our estimates of r, μ and σ^2, suggest that $\lambda' = -4.22$. Because $\lambda' < 0$ and dividends have a positive trend, estimates of $D_t^{\lambda' - 1}$ will be of vanishing importance in explaining prices. Indeed, when we included $D_t^{\lambda' - 1}$ in the regression, it had no effect on the estimate of c_1. Furthermore, c_2 was imprecisely estimated and varied widely across different estimation techniques. As we expected, there seemed to be no evidence that the second nonlinear term helped in explaining stock prices. We therefore do not report these results.

[31] Despite these differences in point estimates, there is virtually no improvement in R^2. A likelihood ratio test cannot reject the hypothesis that line (3) is no improvement over line (1) of Table 3.

[32] The derivative of the likelihood function with respect to the parameters c and λ includes the highly correlated terms $D_t^{\lambda - 1}$ and $c_0(\lambda - 1)D_t^{\lambda - 2}$ (recall that λ is estimated to be greater than 2).

hypothesis is rejected strongly at any reasonable level of significance.[33]

The finding that c is statistically positive suggests that prices become increasingly overvalued relative to the no-bubble price, P_t^{pv}, as dividends rise. Similarly, when dividends are low, the bubble component of price shrinks – P_t approaches P_t^{pv}. The dotted curve in Figure 1 graphs the relationship between fundamentals and prices implied by $c > 0$. The size of the bubble – the distance between P_t and P_t^{pv} – explodes for extreme values of the dividend. Of course, if realized dividends have not spent much time in the explosive range, the bubble component may be quite small.

Note also that the model's estimates of c_0 are sensible. All four estimates from Table 3 imply that P_t^{pv} is measured on average to be about 14 times current dividends; indeed, each estimate is statistically indistinguishable from $\kappa = 14.0$, the value predicted by the present-value model above. In our estimates of (13), $\hat{P}_t^{pv} = \hat{c}_0 D_t$ turns out to be consistent with the long-run average return on stocks because the nonlinear dividend term soaks up a reasonable amount of the excessive sensitivity of actual prices to dividends.

The economic significance of the bubble is, of course, another matter. How large is the bubble component in prices, and how well does the model track actual price movements? Figure 5 helps explore these issues. It compares actual stock prices with the model's estimate of both P_t^{pv} (the no-bubble component of prices) and $\hat{P}_t$ (the model's estimated price inclusive of the bubble terms). Figure 5a presents comparable graphs of price-dividend ratios.[34] The figures are striking in two respects.

First there is the sheer size of the bubble itself – the distance between $\hat{P}_t$ and $\hat{P}_t^{pv}$. It has grown over time and has been particularly large during the post-World War II period. Indeed, the estimates suggest that at this writing the no-bubble level of the Dow-Jones Average is 1,210 – less than 50 percent of its current value! The difference $\hat{P}_t - \hat{P}_t^{pv}$ is estimated to be this large recently because the levels of both dividends and price-dividend

[33] Formally, we should (and will in the next version) include a χ^2 test of this hypothesis. For now, note that setting $c = 0$ in (13) yields an R^2 of 0, against the R^2s reported in Table 3.

[34] Figures 5 and 5a use the estimated coefficients from the third line of Table 3. However, this choice is immaterial to the results: it is almost impossible to distinguish visually among all the models estimated in Table 3.

ratios are historically high.

Second, Figure 5 indicates that $\hat{P}_t$ explains a good deal of actual stock price movements. The sustained runup in prices from 1950 to 1968 appears to be captured by the model, as does the post-World-War-II tendency for stocks to sell at historically large multiples of dividends. The model also does a plausible job of explaining the variability of stock prices. Note from Figure 5 that the variance of dividends appears to have fallen relative to the variance of prices over the sample. Stock-price variability has been somewhat of a puzzle not only because it is so large, but also because it has not declined over time as rapidly as has the variability of dividends. Figure 5 and (13) together suggest a resolution to this paradox: stock price volatility has not fallen with that of dividends because the *level* of dividends – and therefore the scope for volatility due to an intrinsic bubble – has been historically high.[35]

Of course, the "fit" of $\hat{P}_t$ in Figure 5 cannot be judged without a standard of comparison. Because there are an infinite number of bubble specifications which depend on time and/or other extraneous variables, sufficient excavation would allow us in principle to fit perfectly the actual price path. We merely compare our restrictive version of a dividend bubble with the similarly restrictive time-driven bubble $\tilde{P}_t$ defined in (11). Figure 6 graphs the predicted values of the present-value price, $\tilde{P}_t^{pv}$, and the bubble-inclusive price, $\tilde{P}_t$, from OLS estimates of that equation. The parameter estimates are presented in the first two lines of Table 4.

It is immediately clear from Table 4 and Figure 6 that the time-bubble, $\tilde{P}_t - \tilde{P}_t^{pv}$, is neither statistically nor economically very important for understanding stock prices. In addition, the estimates of the constant term, b_0, are less reasonable than those of c_0 in Table 3. Correlation with dividends clearly is not enough to enable the time-driven bubble

[35] To see how much the estimated sensitivity of prices to dividends has changed over time, recall that $dP_t/dD_t = \kappa + c\lambda D^{\lambda-1}$. Using the estimates from Table 3 we can compute rough estimates of dP_t/dD_t, which can be interpreted as the model's prediction of the coefficient in a "cointegrating" regression of prices on dividends. Using average dividends over the period 1951-88 we find (using line 2 of Table 3) $dP_t/dD_t \approx 14.2 + (.26)(2.74)(7.86^{1.74}) = 39.9$. Similarly, over the period 1900-50, $dP_t/dD_t \approx 14.2 + (.26)(2.74)(4.31^{1.74}) = 23.2$. The estimated sensitivity of prices to dividends has therefore nearly doubled over the post-World War II period.

to explain stock prices.

To round out this section, the rest of Table 4 presents estimates of c_0 and c in (13) adding various additional terms: the time-driven bubble term in (14) and a linear time trend. When the price-dividend ratio is regressed on these terms in isolation, they are statistically significant. However, *neither* remains statistically significant when the nonlinear term in (13), $D_t^{\lambda-1}$, is added to the regression. Note that even the sign of the coefficients on the time bubble and linear trend become negative when $D_t^{\lambda-1}$ is added. The coefficient on the nonlinear terms, however, remains statistically significant and essentially unchanged in magnitude.

To see if the nonlinearity in dividends of (13) is important, lines (5) and (6) in Table 4 add a *linear* dividend term, D_t, to the regression. Analogously to the lines above, D_t is positive and statistically significant on its own. But when $D_t^{\lambda-1}$ is included, the sign of the coefficient on D_t is reversed. The sign and magnitude of c is once again not importantly affected.

5. Summary and concluding remarks

This paper has proposed a class of rational bubbles that depend exclusively on exogenous fundamentals. This general class of solutions has intuitive appeal because it does not require the introduction of extraneous variables yet captures the idea that prices can be excessively volatile relative to fundamentals.

We applied a basic version of our model to U.S. stock-market data. The estimates strongly reject the hypothesis that there is no bubble. They also help to reconcile the historical return on stocks with the level of the price-dividend ratio (and with its correlation with dividends), something that present-value models appear unable to do. In addition, the estimates imply that the bubble component in today's stock prices is very large.

The test statistics above have desirable statistical properties because of the tight parametric form of intrinsic bubbles. Unlike general specification tests, our estimates are consistent under both the null and alternative hypothesis.

Our formulation allows variables such as the price-dividend ratio to predict excess returns. To carry out statistical inference we do require that dividends themselves cannot be used to forecast returns, but, in any case, there is little direct evidence to the contrary. By relaxing the present-value assumption, the tests allow the data to allocate deviations from the present-value model across a bubble term and predictable excess returns. Our interpretation of section 4's results is that, once intrinsic bubbles are permitted, the predictability of excess returns no longer appears to be the central cause of the simple present-value model's failure.

Notwithstanding our empirical results, we, too, find the notion of bubbles somewhat problematic. It is difficult to believe that the market is literally stuck for all time on a path along which price-dividend ratios eventually explode. If the market began on such a path, surely investors would at some point attempt the kind of infinite-horizon arbitrage which rules bubbles out in theoretical models; and since *fully* rational agents would anticipate such attempts, bubbles could never get started. It seems to us an empirical question, however, whether this much foresight should be ascribed to the market. Perhaps agents do not really have as clear a picture of the distant future as the simplest rational expectations models suggest. Stock prices and dividends could follow a nonlinear relation such as the one we estimate for some time before market participants catch on to the unreasonable implications of very high dividend realizations.

6. Appendix 1: Time-series properties of dividends

In deriving (13), we assumed that the log dividend process follows a martingale with trend. We examine briefly the time-series evidence on the dividend-generating process to see if it is consistent with this assumption.

We first test to see if the data can reject the hypothesis that the log dividend process, d_t, contain a unit root.[36] We perform the unit-root tests allowing for alternative assumptions about the presence of a time trend. Neither produces significant evidence against the unit-root hypothesis at the 10 percent level. This result suggests that we can estimate consistently the parameters of the dividend process, μ and σ, by an OLS regression of the change in log dividends on a constant. As noted in the text, the data report these as $\mu = 0.011$ and $\sigma = 0.122$.[37]

Of course, if the solution in (7) is to be correct, investors' conditional expectation of d_{t+1} must equal $\mu + d_t$. It follows that the disturbance ξ_{t+1} in (5) must not only be unpredictable given the past history of dividends, it must also be unpredictable given any broader time-t information set. In particular, because investors' forecast of future dividends must depend on current dividends only, stock prices (which presumably reflect information beyond that in dividends) should not help current dividends predict future dividends. This is a strong assumption, so we check to see how well it fares in the data.

Table A1 reports tests for Granger-causality from prices to dividends. In the first line, we regress log dividend changes on a constant, and the lags of both log dividend changes and log price-dividend ratios. Because the price-dividend ratio should include *all* information relevant for forecasting future dividends, it should pick up any forecastable non-trend component of dividend changes. The table reports the sum of the coefficients on $p_t - d_t$ and it lags, as well as an F-test of the hypothesis that these coefficients are

[36] The tests are those proposed by Phillips and Perron (1986) and Phillips (1987). We allow for fourth-order serial correlation in the residuals, as suggested by those authors. For similar tests see Kleidon (1986) and Campbell and Shiller (1987).

[37] There is some evidence that the residuals in this regression are not white, indicating that a more complex ARIMA process might perform better. The Durbin-Watson statistic was 1.65 – which is inconclusive – and a Q-27 test rejects the hypothesis of no serial correlation in the residuals at a 3.8 percent level of significance.

jointly zero. This test shows that we cannot reject the hypothesis that $p_t - d_t$ has no incremental power for forecasting future dividend changes.[38] This formulation, however, is unnecessarily restrictive. If log prices and log dividends are cointegrated of order $(1, 1)$ but with a coefficient other than 1, our inferences may not be valid. In the second line of Table A1, we therefore run a less restrictive regression, which asks directly whether log prices Granger-cause log dividends. Once again, the data do not reject the hypothesis that log prices have no incremental predictive power for future log dividends.[39]

In their tests of the present-value model, Campbell and Shiller (1987) report evidence to the contrary – that the spread does Granger-cause future dividend changes. However, these rejections appear to depend exclusively on a different convention for dating prices and dividends: those authors use the beginning-of-period price, P_{t+1}, and the average of the previous period's dividend, D_t, to predict average period-t+1 dividends, D_{t+1}.[40] If P_{t+1} contains cleaner, more up-to-the-minute information about the beginning-of-period-t+1 dividend than does the time-averaged variable D_t, then one would expect to find Granger-causality using the Campbell-Shiller dating convention, even when stock prices contain no information beyond that in the past history of dividends. Furthermore, as we argued above (footnote 22), substantial information about the current year's dividends could become available during the month of January. There is therefore little basis for rejecting the hypothesis that prices do not Granger-cause dividends. While the view that prices contain information beyond that in current dividends is certainly a plausible one, there just is not much evidence in its favor in these data.[41] We conclude that (5) is a

[38] We also ran this test in levels rather than logs, using what Campbell and Shiller (1987) call the spread, $P_t - \kappa D_t$, in place of the log price-dividend ratio, and ΔD_t in place of Δd_t. The results, using various measures of κ, are not importantly different from those reported above.

[39] Sims, Stock and Watson (1988) and West (1988b) give the asymptotic justification for this procedure. In both regression tests we used a lag length of 4. Similar tests on alternative lag lengths yielded the same results. We also duplicated these tests on the 1871-1986 data set used by Campbell and Shiller (1987), with no change in the results.

[40] Following the dating convention mentioned at the beginning of section 4, we instead use the beginning-of-period-t price, P_t, along with D_t to predict D_{t+1}. Engle and Watson (1985) also use this convention, and obtain Granger-causality results similar to ours.

[41] We ran the regressions reported in Table A1 using Campbell and Shiller's dating convention, and found results similar to theirs. Campbell and Shiller choose their dating convention because – unlike us – they are concerned with predicting prices (in addition to dividends). However, if there is substantial additional information about future dividends in stock prices, then one might nevertheless expect to find that prices Granger-cause dividends using our dating convention. The results in Table A1, however, suggest that this is not the case.

reasonable empirical approximation to the true process generating dividends.

7. Appendix 2: Derivation of the finite-sample distribution of $\hat{c}$ in (13).

Consider the model, $Y_t = c_0 + cX_t + \eta_t$, where $t = 1...T$, $Y_t = P_t/D_t$, c is a parameter to be estimated, $X_t = D_t^{\lambda-1}$, and the log of D_t evolves according to (5):

$$d_t = \mu t + d_0 + \sum_{s=1}^{t} \xi_s. \tag{A1}$$

Data with sample means removed are denoted by y_t and x_t. Let $\tilde{x}$ represent the random sequence of regressors from time 1 to T, a particular realization of which is given by x. We wish to derive the distribution of the test $\hat{c} = c$, where $\hat{c}$ is the OLS estimate of c. To do this we require the following assumptions:

Assumption 1: The residuals, η_t, are normally, but not necessarily identically or independently, distributed with unconditional mean 0 and autocorrelation function $\delta(k)$.

Assumption 2: The dividend innovations, ξ_t, are independently distributed of the residuals, η_t, at all leads and lags, and have mean 0 and variance σ^2.

Note that the OLS estimate of c is given by:

$$\hat{c}(x) = \frac{\sum_{t=1}^{T} \tilde{x}_t \eta_t}{\sum_{t=1}^{T} \tilde{x}_t^2} = \sum_{t=1}^{T} \left(\frac{\tilde{x}_t}{\sum_{s=1}^{T} \tilde{x}_s^2} \right) \eta_t \equiv \sum_{t=1}^{T} \tilde{w}_t \eta_t, \tag{A2}$$

where the $\tilde{w}_t$ are a random set of weights, which by assumption 2 are independently distributed of the η_t's. By assumptions 1 and 2 we have that the linear combination in (A2), for a given sample path of the regressors, x, is a weighted average of normals, and is therefore normally distributed:

$$\hat{c}(x) - c \sim N(0, (x'x)^{-1} x' D x (x'x)^{-1}), \tag{A3}$$

where $D_{i,j} = \delta(i-j)$. Notice that since the distribution of $\hat{c}$ depends on the particular realization, x, the unconditional distribution of $\hat{c}$ will be a mixture of normals, and will therefore have fat tails. Nevertheless, under both the null and alternative hypotheses, c is estimated consistently.

Even though the *unconditional* distribution of $\hat{c}$ is not normal, the t statistic for $\hat{c}(x) = c$ is distributed $N(0,1)$ in finite samples, provided that D is known. To see this note that from (A3) the t statistic is given by:

$$\frac{\hat{c}(x) - c}{\sqrt{(x'x)^{-1}x'Dx(x'x)^{-1}}} \sim N(0,1). \tag{A4}$$

Because this distribution does not depend on the sample realization, x, it holds *unconditionally*. This is true under both the null and alternative hypotheses.

Of course, (A4) assumes that D is known. If D must be estimated, then the expression on the right-hand side of (A4) does not have a t distribution in finite samples (which would be the case if η_t were serially uncorrelated), but will instead converge in distribution as $T \longrightarrow \infty$ to $N(0,1)$.

8. References

Barsky, Robert, and J. Bradford DeLong, "Why Have Stock Prices Fluctuated?," NBER, March 1989.

Black, Fischer, "Noise," *Journal of Finance*, 41 (1986), 529-43.

Blanchard, Olivier J., and Mark W. Watson, "Bubbles, Rational Expectations, and Financial Markets," in *Crises in the Economic and Financial Structure*, edited by Paul Wachtel. Lexington, MA: Lexington Books, 1982.

Campbell, John Y., and Albert S. Kyle, "Smart Money, Noise Trading and Stock Price Behavior," NBER Technical Working Paper no. 71, October 1988.

Campbell, John Y., and Robert J. Shiller, "Cointegration and Tests of Present Value Models," *Journal of Political Economy*, 95 (October 1987), 1062-88.

Campbell, John Y., and Robert J. Shiller, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1 (Fall 1988), 195-228. (a)

Campbell, John Y., and Robert J. Shiller, "Stock Prices, Earnings, and Expected Dividends," *Journal of Finance*, XLIII (July 1988), 661-76. (b)

Casella, Alessandra, "Price Bubbles and the German Hyperinflation," University of California, Berkeley, 1988.

Cecchetti, Stephen G.; Pok-sang Lam; and Nelson C. Mark, "Mean Reversion in Equilibrium Asset Prices," Ohio State University, 1989.

Cumby, Robert E.; John Huizinga; and Maurice Obstfeld, "Two-Step Two-Stage Least Squares Estimation in Models with Rational Expectations," *Journal of Econometrics*, 21 (April 1983), 333-55.

DeLong, J. Bradford; Andrei Shleifer; Lawrence Summers; and Robert Waldman, "The Economic Consequences of Noise Traders," NBER, 1988.

Diba, Behzad T., and Herschel I. Grossman, "Explosive Rational Bubbles in Stock Prices?" *American Economic Review*, 78 (June 1988), 520-530. (a)

Diba, Behzad T., and Herschel I. Grossman, "The Theory of Rational Bubbles in Stock Prices," *Economic Journal*, 98 (September 1988), 746-754. (b)

Durlauf, Steven N., and Robert E. Hall, "Determinants of Noise in the Dividend Based Stock Price Model," Stanford University, 1988.

Engle, Robert F., and Mark W. Watson, "Applications of Kalman Filtering in Econometrics," paper presented at World Congress on the Econometric Society, Cambridge, MA, August 1985.

Flavin, Marjorie A., "Excess Volatility in the Financial Markets: A Reassessment of the Empirical Evidence," *Journal of Political Economy,* 91 (1983), 929-56.

Flood, Robert P., and Peter M. Garber, "Market Fundamentals versus Price Level Bubbles: The First Tests," *Journal of Political Economy,* 88 (August 1980), 745-70.

Flood, Robert P., and Robert J. Hodrick, "Asset Price Volatility, Bubbles and Process Switching," *Journal of Finance,* 41 (September 1986), 831-42.

Flood, Robert P., and Robert J. Hodrick, "Testable Implications of Indeterminacies in Models with Rational Expectations," NBER Working Paper no. 2903, March 1989.

Flood, Robert P.; Robert J. Hodrick; and Paul Kaplan, "An Evaluation of Recent Evidence on Stock Market Bubbles," NBER Working Paper no. 1971, July 1986.

Froot, Kenneth A., "Tests of Excess Forecast Volatility in the Foreign Exchange and Stock Markets," NBER Working Paper no. 2362, March 1988.

Froot, Kenneth A., and Maurice Obstfeld, "Stochastic Process Switching: Some Simple Solutions," NBER Working Paper no. 2998, June 1989.

Gordon, Myron, *The Investment, Financing, and Valuation of the Corporation.* Homewood, IL: Irwin, 1962.

Gourieroux, C.; J.J. Laffont; and A. Monfort, "Rational Expectations in Dynamic Linear Models: An Analysis of the Solutions," *Econometrica,* 50 (1982), 409-25.

Hamilton, James D., and Charles H. Whiteman, "The Observational Implications of Self-Fulfilling Expectations," *Journal of Monetary Economics,* 16 (November 1985), 353-73.

Kleidon, Allan W., "Variance Bounds Tests and Stock Price Valuation Models," *Journal of Political Economy,* 94 (October 1986), 953-1001.

LeRoy, Stephen F., and Richard D. Porter, "The Present-Value Relation: Tests Based on Implied Variance Bounds," *Econometrica,* 49 (1981), 555-74.

Lucas, Robert E., Jr., "Asset Prices in an Exchange Economy," *Econometrica,* 46 (November 1978), 1426-46.

Mankiw, N. Gregory; David Romer; and Matthew Shapiro, "An Unbiased Reexamination of Stock Market Volatility," *Journal of Finance*, 40 (1985), 677-87.

Marsh, Terry, and Robert Merton, "Dividend Variability and Variance Bounds Tests for the Rationality of Stock Market Prices," *American Economic Review*, 76 (1986), 483-98.

McCallum, Bennett T., "On Non-Uniqueness in Rational Expectations Models: An Attempt at Perspective," *Journal of Monetary Economics*, 11 (March 1983), 139-68.

Newey, Whitney, and Kenneth West, "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55 (May 1987), 703-8.

Phillips, Peter C.B., "Time Series Regression with Unit Roots," *Econometrica*, 55 (March 1987), 277-302.

Phillips, Peter C.B., and Samuel Ouliaris, "Testing for Cointegration," Discussion Paper no. 809, Yale University, Cowles Foundation, October 1986.

Phillips, Peter C.B., and Pierre Perron, "Testing for Unit Roots in Time Series Regression," Yale University, Cowles Foundation, June 1986.

Pindyck, Robert S., "Risk, Inflation, and the Stock Market," *American Economic Review*, 74 (1984), 335-51.

Poterba, James M., and Lawrence H. Summers, "The Persistence and Volatility of Stock Market Fluctuations," *American Economic Review*, 76 (1986), 1142-51.

Saracoglu, Rusdu, and Thomas J. Sargent, "Seasonality and Portfolio Balance under Rational Expectations," *Journal of Monetary Economics*, 4 (August 1978), 435-58.

Sargent, Thomas J., *Macroeconomic Theory*. Second edition. Orlando, FL: Academic Press, 1987.

Shiller, Robert J., "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?," *American Economic Review*, 71 (1981), 421-36.

Shiller, Robert J., "Stock Prices and Social Dynamics," *Brookings Papers on Economic Activity* (2: 1984), 457-98.

Sims, Christopher A.; James H. Stock; and Mark W. Watson, "Inference in Linear Time Series Models with Some Unit Roots," Stanford University, February 1988.

Summers, Lawrence H., "Does the Stock Market Rationally Reflect Fundamental Values?," *Journal of Finance,* 41 (July 1986), 591-601.

West, Kenneth D., "A Specification Test for Speculative Bubbles," *Quarterly Journal of Economics,* 102 (August 1987), 553-80.

West, Kenneth D., "Bubbles, Fads, and Stock Price Volatility Tests: A Partial Evaluation," *Journal of Finance,* 43 (July 1988), 639-55. (a)

West, Kenneth D., "Asymptotic Normality when Regressors Have a Unit Root," *Econometrica,* 56 (November 1988), 1397-1418. (b)

Whiteman, Charles H., *Linear Rational Expectations Models: A User's Guide.* Minneapolis, MN: University of Minnesota Press, 1983.

Figure 1
Intrinsic-Bubble Price Paths

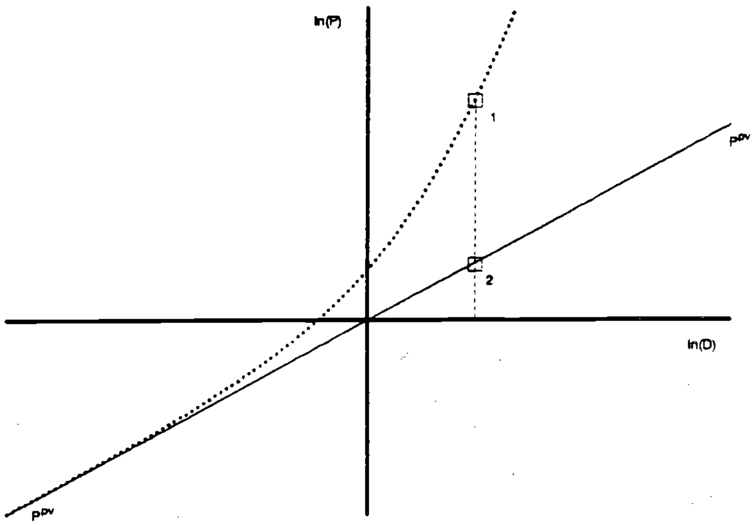

Figure 2
Simulated Stock Price Paths

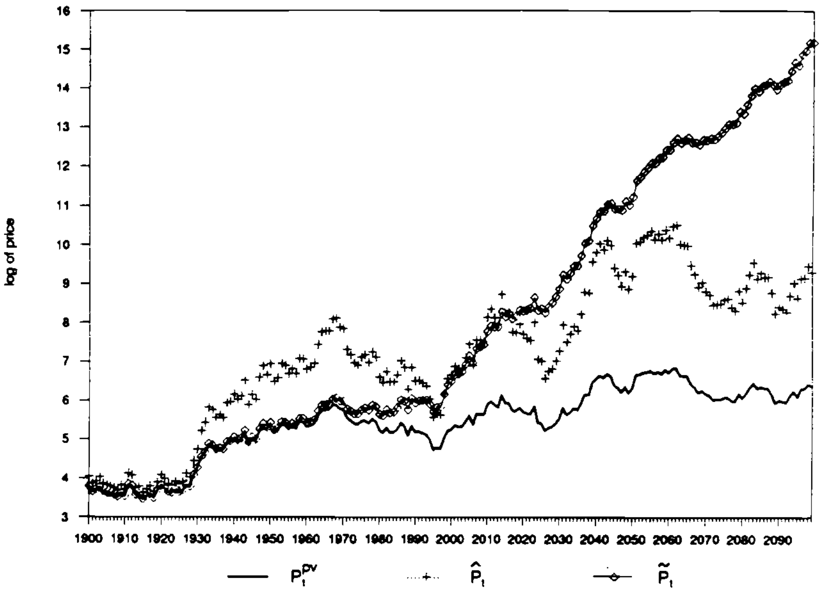

Figure 3
Simulated Stock Price Paths

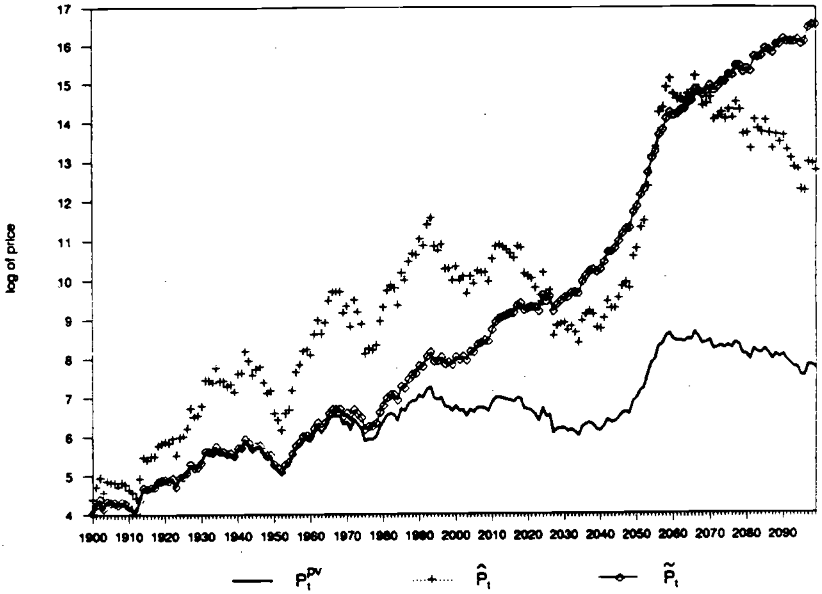

Figure 4
Simulated Stock Price Paths

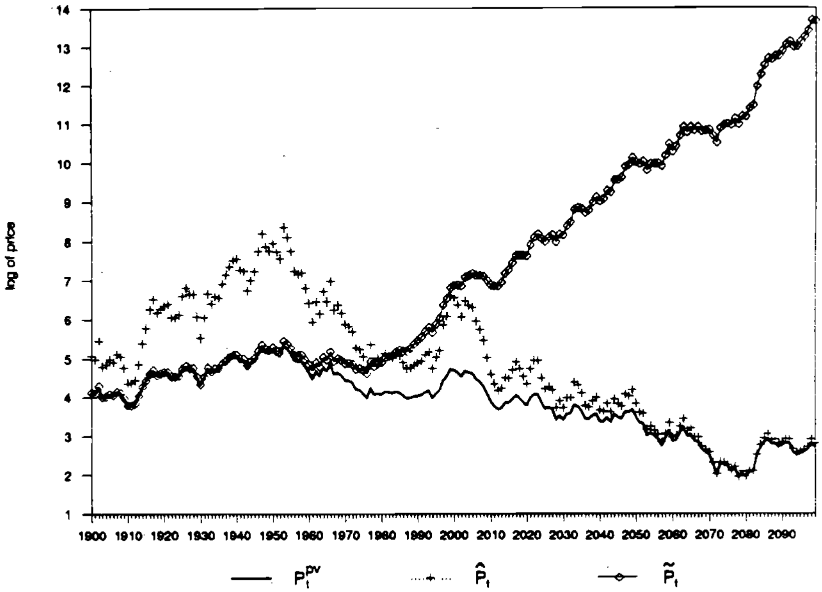

Table 1
Cointegrating Regressions

Row / Regression Equation		Cointegrating Coefficient	R^2	DW	DF
(1)	$P_t = \alpha + \beta D_t$	36.65 (1.63)	0.85	0.57	87
(2)	$D_t = \alpha + \beta P_t$	0.023 (0.001)	0.85	0.69	87
(3)	$p_t = \alpha + \beta d_t$	1.591 (0.06)	0.88	0.69	87
(4)	$d_t = \alpha + \beta p_t$	0.556 (0.020)	0.88	0.70	87

Notes to Table 1: Cointegrating regressions are estimated using OLS, with uncorrected standard errors in parentheses. Sample period for all regressions was 1900-88.

Table 2
Unit Root Tests

Row / Series	with time trend	without time trend
(1) Log Dividends	-0.1644	-0.0545
	-2.56	-1.40
(2) Price-Dividend Spread, P - 14.0 D	-0.1355	-0.0702
	-2.08	-1.14
(3) Price-Dividend Ratio, P / D	-0.2157	-0.1343
	-2.99	-2.11
(4) Log Price-Dividend Ratio, p - d	-0.2122	-0.1315
	-3.55 **	-2.55

Notes to Table 2: Figures reported are the coefficients β_1 in the following regressions: with trend, $\Delta x_{t+1} = \beta_0 + \beta_1 x_t + \beta_2 t$; and without trend, $\Delta x_{t+1} = \beta_0 + \beta_1 x_t$. Standard errors are constructed allowing for an MA(4) process in the residuals. T-statistics reported beneath the point estimates are for the test $\beta_1 = 0$. *, **, *** represent significance at the ten, five, and one percent levels, respectively, using confidence intervals proposed by Phillips and Perron (1986) and Perron (1987).

Table 3
Estimates of

$$P_t/D_t = c_0 + cD_t^{\lambda-1}$$

Row / Estimation method	c0	c	λ-1	F-test c =0	R^2	DW	DF
1. OLS	12.24	0.34			0.57	0.71	87
	(1.14)	(0.05)***					
2. Maximum Likelihood	14.18	0.26			0.75	1.91	86
	(1.77)	(0.06)***					
2. OLS	14.63	0.04	2.61	128.0 ***	0.57	0.71	86
	(2.28)	(0.12)	(1.15)**				
5. Maximum Likelihood	16.55	0.01	3.29	9.62 ***	0.75	1.92	85
	(2.02)	(0.02)	(1.45)**				

Notes to Table 3: OLS regressions report Newey-West standard errors, allowing for fourth-order serial conditional and heteroskedasticity. (Higher orders of serial correlation did not yield larger standard errors.) Maximum likelihood estimates specify the error term as an AR(1) process. *, **, *** represent statistical significance at the ten, five, and one percent levels, respectively. Sample period for all regressions was 1900-88.

Figure 5
Actual and Predicted Stock Prices

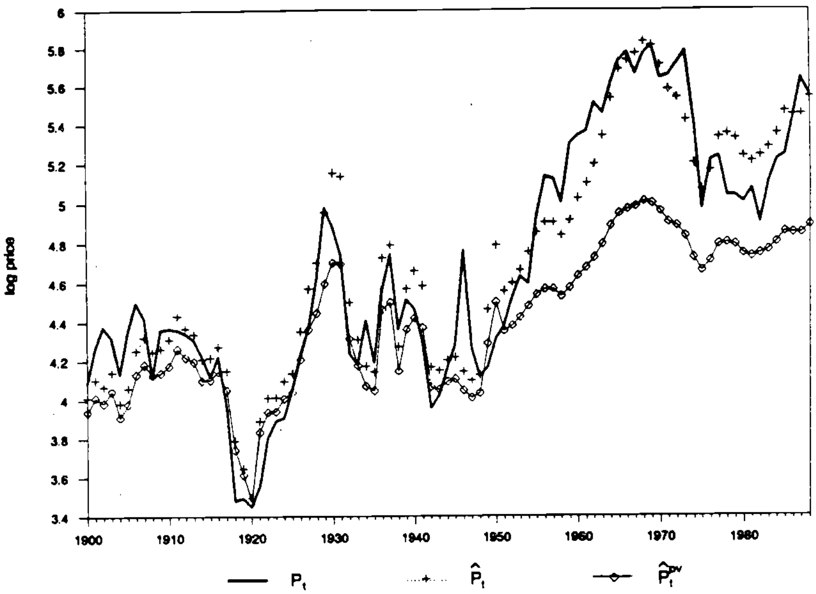

Figure 5a

Actual and Predicted Log Price-Dividend Ratios

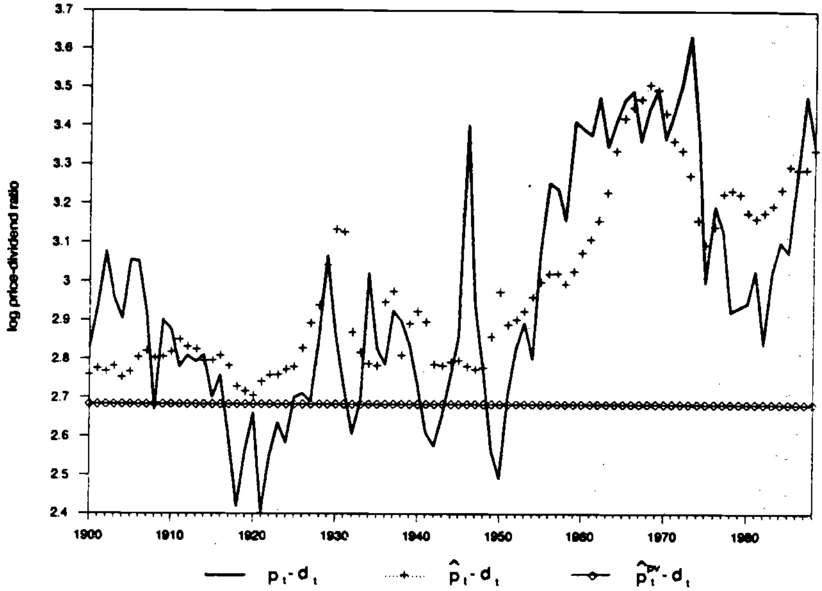

Figure 6
Actual and Predicted Stock Prices

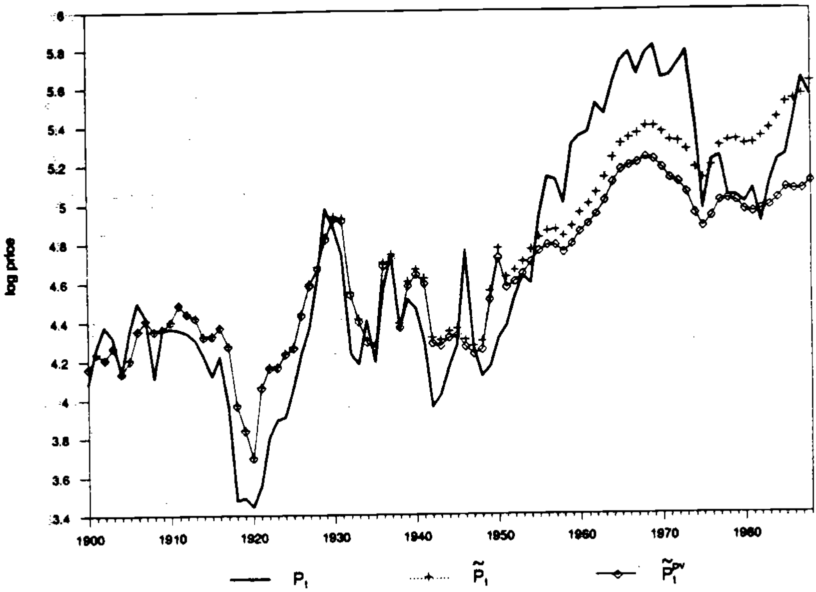

Table 4
Estimates of Alternative Models

$$P_t/D_t = c_0 + cD_t^{\lambda-1} + gX_t$$

Row / Estimation method	c0	c	g	R^2	DW	DF
1. OLS, time bubble	18.28		0.030	0.21	0.35	87
	(1.47)		(0.013)**			
2. OLS, time bubble	11.85	0.377	-0.008	0.57	0.75	86
	(1.09)	(0.060)***	(0.010)			
3. OLS, linear trend	13.68		0.15	0.35	0.43	87
	(2.22)		(0.05)***			
4. OLS, linear trend	12.43	0.364	-0.019	0.57	0.75	86
	(1.39)	(0.066)***	(0.052)			
5. OLS, linear dividends	6.88		2.273	0.55	0.70	87
	(2.02)		(0.39)***			
6. OLS, linear dividends	18.09	0.684	2.397	0.57	0.70	86
	(8.62)	(0.448)	(3.32)			
1. Maximum Likelihood, time bubble	18.34		0.027	0.75	2.04	86
	(2.16)		(0.014)*			
2. Maximum Likelihood, time bubble	14.48	0.223	0.008	0.75	1.94	85
	(1.86)	(0.075)***	(0.012)			
3. Maximum Likelihood, linear trend	14.02		0.145	0.75	1.99	86
	(2.87)		(0.054)***			
4. Maximum Likelihood, linear trend	13.19	0.190	0.060	0.75	1.93	85
	(2.16)	(0.084)**	(0.056)			
5. Maximum Likelihood, linear dividends	11.39		1.530	0.75	1.93	86
	(2.76)		(0.420)***			
6. Maximum Likelihood, linear dividends	24.54	0.904	-4.372	0.76	1.91	85
	(6.64)	(0.404)**	(2.710)			

Notes to Table 4: The variable X_t is given by: a time bubble ($X_t = e^{(r-\mu-\sigma^2)t}$); a linear trend ($X_t = t$); and linear dividends ($X_t = D_t$). OLS regressions report Newey-West standard errors, allowing for fourth-order serial conditional and heteroskedasticity. (Higher orders of serial correlation did not yield larger standard errors.) Maximum likelihood estimates specify the error term as an AR(1) process. *, **, *** represent statistical significance at the ten, five, and one percent levels, respectively. Sample period for all regressions was 1900-88.

Table A1

Tests for Whether Prices Granger-cause Dividends

Row / Regression Equation		F-test (p-value)	R^2	DW	DF	lag length
(1)	$\Delta d_{t+1} = \alpha(L)\Delta d_t + \beta(L)(p_t - d_t)$	0.812 (0.52)	0.13	1.96	75	4
(2)	$d_{t+1} = \alpha(L)d_t + \beta(L)p_t$	1.868 (0.12)	0.91	1.98	75	4

Notes to Table A1: Granger causality tests are estimated using OLS. The sum of the coefficients on the log price-dividend ratio in line (1) and on the log of price in line (2) are reported. In the parentheses below these are probability values from F-tests of the hypothesis that $\beta_i = 0, \ \forall i$. Alternative lag lengths were also tried for these regressions, but did not change the results. The sample period for all regressions is 1900-88.