



What Happened to the U.S. Stock Market? Accounting for the Past 50 Years

Michele Boldrin and Adrian Peralta-Alva

The extreme volatility of stock market values has been the subject of a large body of literature. Previous research focused on the short run because of a widespread belief that in the long run the market reverts to well-established fundamentals. The authors' research suggests this belief should be questioned. First, they show actual dividends cannot account for the secular trends of stock market values. They then consider a more comprehensive measure of capital income, which displays large secular fluctuations that roughly coincide with changes in stock market trends. Under perfect foresight, however, this measure fails to properly account for stock market movements. The authors thus abandon the perfect foresight assumption and instead assume that forecasts of future capital income are performed using a distributed lag equation and information available up to the forecasting period only. They find that standard asset-pricing theory can be reconciled with the secular trends in the stock market. This study, nevertheless, leaves open an important puzzle for asset-pricing theory: The market value of U.S. corporations was much lower than the replacement cost of corporate tangible assets from the mid-1970s to the mid-1980s. (JEL E25, G12)

Federal Reserve Bank of St. Louis *Review*, November/December 2009, 91(6), pp. 627-46.

Standard (consumption-based) asset-pricing models have a hard time explaining *high-frequency* fluctuations in stock market values, given observed fluctuations in fundamentals. The anomalies these models face have been labeled in a variety of ways—all ending with the word “puzzle.” Various solutions have been suggested, but it seems no more than a handful of researchers have deemed any of these satisfactory. Campbell (2003) provides a summary of the various puzzles and solutions proposed by the consumption-based asset-pricing literature, and Boldrin, Christiano, and Fisher (2001) offer the solution that, so far, we find the least unconvincing. While the search continues, it becomes more apparent that the hope of capturing the stock market’s short-run

gyrations by appropriately filtering the quarterly movements of aggregate consumption is unlikely to be realized. The question then is this: If consumption-based models cannot do the job, what can? In this article we contribute our two cents to the collective effort of answering this most difficult question.

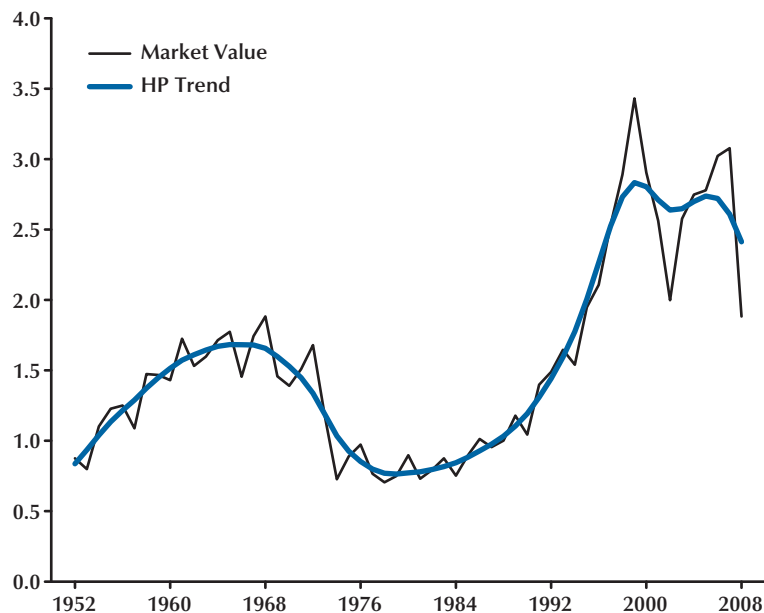
Before completely abandoning the standard model, we ask a seemingly less challenging but more fundamental question. Namely, can the standard (that is to say, net present-value-based) economic theory of asset pricing account for the very large *low-frequency* fluctuations in the aggregate stock market valuation of U.S. corporations? In other words, if we ignore short-term movements and look only at very long-run trends—those persistent enough to last at least five, and generally

Michele Boldrin is the Joseph G. Hoyt Distinguished University Professor in the Department of Economics at Washington University in St. Louis and a research fellow at the Federal Reserve Bank of St. Louis. Adrian Peralta-Alva is an economist at the Federal Reserve Bank of St. Louis. Chanont Banterghansa provided research assistance.

© 2009, The Federal Reserve Bank of St. Louis. The views expressed in this article are those of the author(s) and do not necessarily reflect the views of the Federal Reserve System, the Board of Governors, or the regional Federal Reserve Banks. Articles may be reprinted, reproduced, published, distributed, displayed, and transmitted in their entirety if copyright notice, author name(s), and full citation are included. Abstracts, synopses, and other derivative works may be made only with prior written permission of the Federal Reserve Bank of St. Louis.

Figure 1

Market Value of U.S. Corporations (Relative to Corporate GDP)



more, years—is the standard model capable of correctly explaining/predicting those trends? As far as we know, this question has seldom, if ever, been addressed systematically. Yet, it is relevant for an assessment of asset-pricing models for this reason: There is a widespread belief that, although the standard model may miss a few short-term bumps, in the long run the market will revert to well-established fundamentals. Our research suggests this belief should be questioned.

Our analysis is structured as follows. First, we document the secular trends in the value of U.S. corporations. Available data rule out the possibility that fluctuations in market value might have been caused by fluctuations in corporate assets. Then, we study the implications of a fundamental asset-pricing equation according to which asset prices equal the expected discounted present value of returns. As is common in the literature, to test the implications of the theory, we use aggregate data (either from publicly traded firms or from the overall corporate sector). First, we show that, under perfect foresight, the stan-


dard approach of computing the present value of actual stock market dividends or returns accounts for little of the secular trends of the U.S. stock market. Since dividend payments may respond to complicated corporate finance considerations, we then study whether movements in the whole of shareholders’ income may better account for stock market trends. We find that this is not the case. Finally, we drop the perfect foresight assumption and assume that shareholders make forecasts based only on available information and a distributed lag equation. As we show, this assumption together with the fundamental asset-pricing equation explains much of the secular trends of the U.S. stock market.

THE SECULAR TRENDS OF THE VALUE OF CORPORATE CAPITAL

Figure 1 summarizes key features of the data. Our data appendix details the sources and methods used to construct this and all other figures in this paper. To normalize for economic growth, we

Figure 2

Market Value and Replacement Value of Corporate Capital (Relative to Corporate GDP)



focus on the ratio of the market value of corporations to corporate value added (or other measures of aggregate output, when appropriate); we refer to this as the “market ratio.” The black line in Figure 1 shows the market ratio during the past 55 years (based on annual data). The blue line shows the low-frequency movements in the market ratio by means of the trend generated by the Hodrick-Prescott (HP) filter. The use of the HP filter is standard in the literature and it is applied to all variables in this paper. Almost undistinguishable patterns result when other reasonable long-run filters are used.

As the figure shows, after two decades of growth the market ratio declined by 50 percent during 1973-74 and stagnated until the mid-1980s. From 1985 to 2000 it more than tripled, only to collapse again by 2001. Since then, the market ratio has fluctuated around 2.7, taking a gigantic drop (only partially reported in the figure and now partially recovered) during recent months. These are extremely large movements by any metric and dwarf the also substantial oscillations

at the quarterly to yearly frequencies. The question is this: What kind of economic rationale drives such impressive swings?

The most elementary model we can think of analyzes aggregate production and capital accumulation over time. This type of model considers an aggregate firm producing national consumption (in fact, gross national product [GNP]) by employing capital, K , and labor, L , under a constant returns-to-scale production function, $F(K,L)$. The resource and wealth constraints for this economy are

$$F(K_t, L_t) = C_t + I_t; K_{t+1} = (1 - \mu)K_t + I_t.$$

In this environment, consumption and investment (and therefore capital) are interchangeable on a one-to-one basis. Hence, the price (or value) of the capital stock, measured in units of the consumption good, is always one. It follows that the market ratio must equal the physical capital/output ratio implied by the aggregate production function $F(K,L)$. This most elementary explanation is immediately ruled out by the data. In Figure 2,

we add a third line to Figure 1, which shows the ratio of the replacement value of corporate capital to corporate gross domestic product (GDP) (that is to say, K_t/Y_t). It shows a remarkable stability compared with the market ratio: Although some long-run oscillations are visible, they are of about one order of magnitude smaller than those of the market ratio and in the opposite direction. In summary, an explanation for the huge swings in the market ratio needs to be found somewhere beyond the actual stock of capital owned by U.S. corporations or the cost of producing it.¹

PERFECT FORESIGHT OF FUTURE DIVIDENDS

If oscillations in the market value of capital cannot be explained in terms of either its cost or its quantity (relative to labor and/or output), maybe they can be explained in terms of “value.” The market ratio increases (decreases) because the capital stock used by corporations becomes less (more) productive, hence yielding more (less) profits to its owners. According to this principle, the market value of corporate capital is determined by looking forward and not backward: Independent of both how capital intensive production processes may be and the cost of installed machines, the market value will increase if capital is productive and its owners expect it to yield

¹ A slightly more sophisticated version of this model, which allows for changes in the relative price of capital (q), can be evaluated but yields similar (counterfactual) results. Suppose the investment good is produced with a different technology from the one producing the consumption good (see, e.g., Boldrin, Christiano, and Fischer, 2001). In this case, the resource and wealth constraints read as

$$F(K_t^c, L_t^c) = C_t; G(K_t^k, L_t^k) = I_t; \\ Y_t = C_t + q_t I_t; K_{t+1} = (1 - \mu)K_t + I_t;$$

and the market ratio equals

$$\frac{q_t K_t}{F(K_t^c, L_t^c) + q_t G(K_t^k, L_t^k)}.$$

Hence, either “biased” technological change or changes in sectoral factor intensities could bring about a change in the relative price of capital. Moreover, the market ratio may move because the capital intensity of aggregate production moves or because the relative price of capital oscillates. Notice, however, that the market ratio predicted by the model, expression (2), should still correspond to the capital output ratio of the United States (just as the simpler one-sector model).

profit; and it will decrease in the opposite case. In summary: Standard asset-pricing theory says that the market ratio is a forecast. The questions are (i) the market ratio is a forecast of what and, (ii) how correct this forecast has turned out to be? We will concern ourselves with these two questions for the rest of this paper.

To begin answering them, we establish the simplest possible framework of analysis in which the value of corporations is equal to the value of what their capital will produce and earn. We contemplate dynamic stochastic economies that are, on a period-by-period basis, subject to a vector of exogenously given shocks. Such shocks—which may include changes in productivity, demand, taxes, and other factors—are the sources of uncertainty that the forward-looking agents must examine to price assets on the basis of their expected future returns. Let E_t be the expectation operator, taken with respect to the probability distribution of future shocks on the basis of the information available at time $t = 0, 1, 2, \dots$; d_t are the dividends paid by the firm to shareholders, τ_t^d and τ_t^v the dividends and capital gains income tax rates, and V_t the market value of the firm, all as of period t . Finally, let p_{t+i} be the stochastic discount factor of future consumption—that is, the value today of one unit of consumption obtained, in some state of the world, during the future period $t+i$, $i = 1, 2, 3, \dots$. The following relation holds:


$$(1) \quad V_t = (1 - \tau_t^d)d_t + E_t \left[\sum_{i=1}^T (1 - \tau_{t+i}^d) p_{t+i} d_{t+i} \right] \\ + E_t \left[(1 - \tau_{t+T}^v) p_{t+T} V_{t+T} \right],$$

Here, T denotes any arbitrary positive number (including plus infinity). This formula states that the market value of a firm should equal the expected present discounted value of the future stream of (after-tax) shareholders’ income it generates plus the (after-tax) capital gains/losses that result from selling the share at some future period.

It is important to note that the fundamental asset-pricing equation (1) holds in a wide range of economic models. Indeed, different branches of the literature have emerged from varying the key assumptions and methods for deriving predictions from this equation. The consumption-

Figure 3

Present Value of Dividends versus Market and Replacement Values (Relative to Corporate GDP)



based asset-pricing literature, for instance, assumes dividends and consumption are exogenously given processes. In this literature, the interaction of the stochastic discount factor and the dividend process are the key forces driving the volatility of asset prices. The production-based asset-pricing literature, in turn, develops the asset-pricing implications of models where consumption and dividends are endogenously determined. Finally, the present-value pricing literature considers long holding periods for shares (high values for T , in our notation) and explores the asset-pricing implications of alternative measurements for dividends and long-run discount factors.

To derive the exact quantitative implications of the asset-pricing equation (1), one would need to measure all of the possible time series of taxes, discount factors, capital income, and, in particular, the probabilities the market assigns to all possible future states of the world. Determining the latter directly, at any given point in time, is an impossible task because the theory, per se, admits the most arbitrary set of expectations for the par-

ticipating agents. A common benchmark followed in the literature (e.g., Shiller, 2005) consists of assuming a constant discount factor and approximating capital income by actual dividends. Crucially, perfect foresight on dividends is commonly assumed as well: The market prices are supposed to be based on exact forecasts of the realized dividends; hence, realized dividends can be used in the computations. Because these are open-ended models, existing analyses typically complement the perfect foresight hypothesis with the additional assumption that dividends will grow at some average rate for the infinite future. Generally, this literature also abstracts away from fluctuations in the tax on dividend income.

This approach does not go very far in accounting for equity price movements. A representative illustration of the predictions from the theory under the aforementioned assumptions, based on the dividend data compiled by Shiller (2005), is presented in Figure 3. For comparison, the ratios displayed below have been normalized so that their 1960-72 average is equal to one. (We will

follow this normalization procedure throughout this paper.) These computations assume a 7 percent discount rate and a terminal growth rate for dividends of 3 percent for the infinite future following period t . We use a constant discount rate to simplify our presentation, but our results do not change much if we use instead a discount rate based on a power function of consumption growth, as is standard in consumption-based asset-pricing theory. Our assumption of 3 percent terminal growth for dividends is also incorrect since the dividend-to-GDP ratio has been decreasing over time, while GDP has been growing at an average of about 3 percent for most of the period. The perfect foresight assumption implies that the market ratio should have been very high earlier on—as dividends were a high percentage of corporate GDP in the 1950s and discounting matters—to then subsequently decrease and remain quite stable (as dividends' growth rates stabilized from the mid-1970s onward). This makes the large oscillations that took place between 1970 and today impossible to explain on the basis of dividend payoffs, perfect foresight, and a stable dividends-to-GDP ratio. Our first conclusion is that one, or more, of these three assumptions—dividends are the payoff to be forecasted, they are forecasted correctly, and they grow at roughly the same rate as corporate GDP—must be disposed of.

The fact that equity prices increased so rapidly during the late 1990s and that the value of dividends did not has been interpreted by some as evidence of irrational exuberance. This is not necessarily correct: There may have been “exuberance,” but it need not have been “irrational” insofar as the information available to the market did not need to be of sufficient quality or amount to correctly forecast future dividends. We will return to this point later, as the issue of what the market can and cannot “forecast correctly” is at the root of the problem we are addressing. In any case, an “irrational bubble” might have been partially behind equity prices during the mid-1990s, but notice that even after both the 2000 and recent stock market crashes, the market ratio is much higher than during the 1980s: So, what is it that the market ratio is therefore “forecasting”? Similarly, the issue of why equity prices were so low

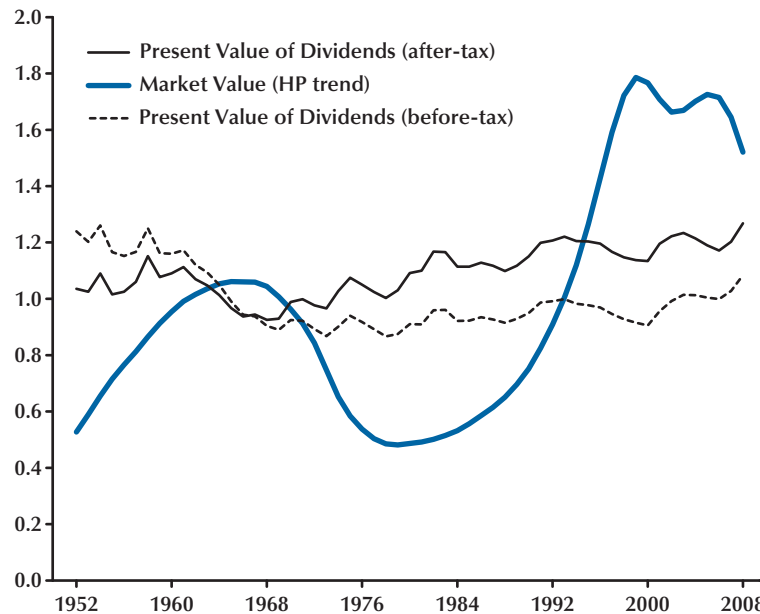
in the early 1950s and in the mid-1970s and -1980s is not frequently addressed in the financial literature either, which again brings up a similar question: What was the market ratio “forecasting” during those periods? Not dividends—or, at least, not correctly—because the net present value of actual future dividends is above the market ratio between the early 1970s and the early 1990s.

As we will argue later, the low market valuations in the middle period appear to be the hardest to understand. Notice in passing that this hypothetical market ratio, computed only on the basis of observed dividends, is much closer to the replacement value of capital than the actual market ratio. In other words, if the stock market had really valued corporations on the basis of actually realized dividends, the market and replacement values of corporate capital would have been relatively close during the 56 years we study, and the only long-run puzzle would be a persistent difference, in levels, between the replacement cost of capital and the present value of the dividends it has been generating. Such a puzzle could be easily solved, though, by lowering the discount rate below the 7 percent value we used in the reported calculation. But, apparently, this is not what the stock market did.

We now refine this “perfect foresight” methodology by modifying the object supposedly forecasted by the market ratio. First, McGrattan and Prescott (2005) document important changes in the taxation of dividend income and investment subsidies that may account for some of the observed fluctuations in equity values. We therefore recompute the implications of the theory by adjusting for the varying rate of dividends' taxation. The results are shown in Figure 4, and they are not good. Which leads us to repeat the observation made earlier: The perfect foresight assumption, when applied to the valuation of future payment streams, imposes strong restrictions on the model's predictions. In particular, by eliminating any learning process and assuming the market “knows” future events much earlier than they materialize, it tends to “front-load” all historical changes, producing (thanks to discounting) very flat predictions. In summary, if agents can more or less perfectly forecast all relevant vari-

Figure 4

Present Value of Dividends (Before- and After-Tax) versus Market Value (Relative to Corporate GDP)



ables, the long-run swings of the market ratio make no sense whatsoever. This suggests that the problem may not be with “what” the market forecasts but with “how” it does it.

It may be possible to offer arguments in favor of the perfect foresight assumption for economic fundamentals (such as dividends and discount factors). Assuming perfect foresight in policy variables such as taxes seems much harder to do. Indeed, McGrattan and Prescott’s (2005) analysis studies the impact of an unexpected and permanent change in taxes.

We thus recompute the predictions for the theory under the assumption of perfect foresight on dividends and interest rates, but assuming that, at each period, a new tax rate arrives unexpectedly and this rate is believed to persist into the infinite future. As first noted by Bian (2007), such changes in dividend income taxes can only (very) partially account for the higher values of the market ratio from 1994 to 2008. However, in this case, the size of the increase is too small and

its timing is way off. Figure 5 suggests that the stock market undervalued corporations between 1952 and 1961 and then, again, between 1970 and 1992, while some kind of “exuberance” (irrational or not, we will see) has driven the market ratio from about 1996 to the present. Simply, even after allowing for large tax surprises, the net present value of future dividends provides a very poor explanation of what happened to the market ratio.

Symmetry would require assuming that a new (permanent) growth rate for dividends also arrives, unexpectedly, at each period. Under these conditions—namely, a random walk growth rate for dividends, g , and a constant discount factor, r —the asset-pricing equation will simply predict that the market ratio moves proportionally with the dividend growth rate. In fact, the random walk hypothesis implies that the market value should equal today’s (after-tax) dividends divided by $r_t - g_t$. Figure 6 shows the growth rate of real dividends and averages (of different lengths) of past growth rates. Notice, first, that the growth rate of

Figure 5

Present Value of Dividends (Before and After Unexpected Tax Changes) versus Market Value (Relative to Corporate GDP)





Figure 6

Growth Rate of Dividends (Actual and Averages of Past Years)



dividends is fairly volatile. Second, up to the mid-1980s, changes in the dividend growth rate roughly coincide with changes in the trend of the market ratio. The growth rate of dividends goes down around 1968, and so does the market ratio. Similarly, dividend growth is low through the mid-1970s and it does not recover until the mid-1980s. The market ratio follows similar patterns. Dividend growth does not have any specific trend, on average, during 1992-2008, but it displays higher volatility. While the big increases in the market ratio of the mid-1990s and later are hard to be accounted for by trends in dividend growth, it is surprising how well the 5- and 10-year averages of the dividends' growth rate mimic the long-run gyrations of the market ratio.

It is important to stress that, although changes in the (average trend of the) dividend growth rate are positively related to changes in the trend of the market ratio, the growth rates of actual dividends are very often negative. Of course, assuming that dividends will grow at a negative rate for the infinite future is not very realistic, which makes us return to our fundamental question. If it is not a forecast of actual dividends paid, then the market ratio forecast is a forecast of what? The sections that follow refine the production-based asset-pricing model to provide one possible answer.

PERFECT FORESIGHT OF TOTAL CAPITAL INCOME

Actual dividends paid, in light of our previous analysis, cannot help explain any of the big historical swings in the market ratio. It is not clear, however, that one should use actual dividends in testing the theory. In particular, actual dividend payments may respond to additional considerations such as informational asymmetries, principal-agent revelation mechanisms, fiscal incentives other than those captured by the taxation of dividends and capital gains, and so on. (Easterbrook, 1984, and Feldstein and Green, 1983, review some of the relevant literature.) In trying to determine how far the fundamental asset-pricing equation (1) can take us, it seems more appropriate to

abstract from dividend payment considerations and start instead from a simple framework whereby firms' net worth equals the present value of *all* shareholders' income.

The earlier model of production can be adapted to this end by assuming that the aggregate firm chooses capital, labor, and investment to maximize the net present value of shareholders' income. In this new version of the model, shareholders' income is endogenously determined. Further, shareholders are the residual claimants of corporate value added after compensation of employees, corporate income taxes, and gross investment are taken care of. It follows that income accruing to shareholders in period t is

$$d_t = Y_t - w_t L_t - I_t - \text{taxes},$$

where w_t is the wage rate and "taxes" includes all kind of taxes that firms are in effect responsible for paying.


What are the asset-pricing implications of this type of model? First, the fundamental equation (1) is still valid. However, we now have a definition for capital income, consistent with a specific theory, that can be easily mapped into the U.S. national income and product account (NIPA) data.² The second prediction of the model is, as before, the familiar identity of market value and the value of all of the firm's assets (capital stock), after adjusting for dividend income taxes:

$$V_t = (1 - \tau_t^d) d_t + E_t \left[(1 - \tau_{t+1}^v) p_{t+1} V_{t+1} \right].$$

This makes it clear that shareholders may obtain income from ownership of the firm in two different ways. The first is dividend payments, which is what the firm supposedly maximizes and that accrues to owners holding shares in perpetuity. The second way to obtain income is by selling equity shares, which may result in capital gains (or losses). Notice, finally, that standard measures of capital income equal shareholder income plus investment expenditures. This is consistent with our model. Investment expendi-

² Company-originated quarterly earnings are more comprehensive measures of capital income than dividends. However, we find that the results we report, when based on earnings, are very similar to those based on actual dividends.

Figure 7
Growth of Corporate GDP



tures are indeed a form of capital income since they may ultimately affect future dividends and the future value of the firm, which are both taken into account by the asset-pricing equation above. Put differently, shareholder total income in period t is the sum of the dividends received and of the (potential) capital gains accrued; the latter include (among other things) the value of period t 's gross investment.

Because shareholder income is, by an accounting identity, equal to the fraction of corporate value added accruing to shareholders multiplied by corporate GDP, it is worth considering the two components separately to see whether their movements over time teach us anything useful. In standard macroeconomic models attention is focused, more often than not, on the time-series behavior of corporate GDP, while the share accruing to the owners of capital is taken as constant and paid little attention. This analytical choice is unfortunate since, as we will show later on, it may lead one to miss a fundamental factor affecting movements in stock market evaluations. Figures 7 and 8

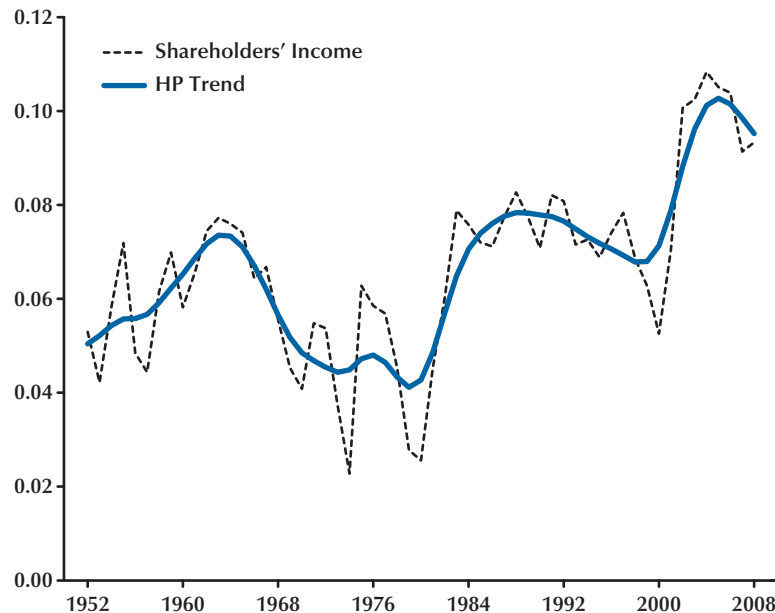
report the two components separately. Notice that these are corporate GDP data and thus do not include the impact of changes in personal taxes in the net income accruing to shareholders.

There are various salient features in these data. First, corporate GDP growth fluctuates widely around an otherwise apparently stable long-run growth rate (with, possibly, a very modest downward trend in the latter period) and there are, really, only two decades of “major” (i.e., above-average) growth: the 1960s and the 1990s. Because these are also the two periods during which the market ratio rallied the most, one would be inclined to say that “roughly” the stock market captured the underlying long-run oscillations in payoff. The key word here, though, is “roughly”; in fact, very roughly as the subsequent quantitative analysis will show. Further, the recovery that ended almost two years ago was nothing spectacular: In terms of total corporate GDP growth, it was the *worst expansion* of the past 50 years!

Second, shareholder income as a fraction of total corporate value added (or capital income

Figure 8

Shareholders' Income (Relative to Corporate GDP)



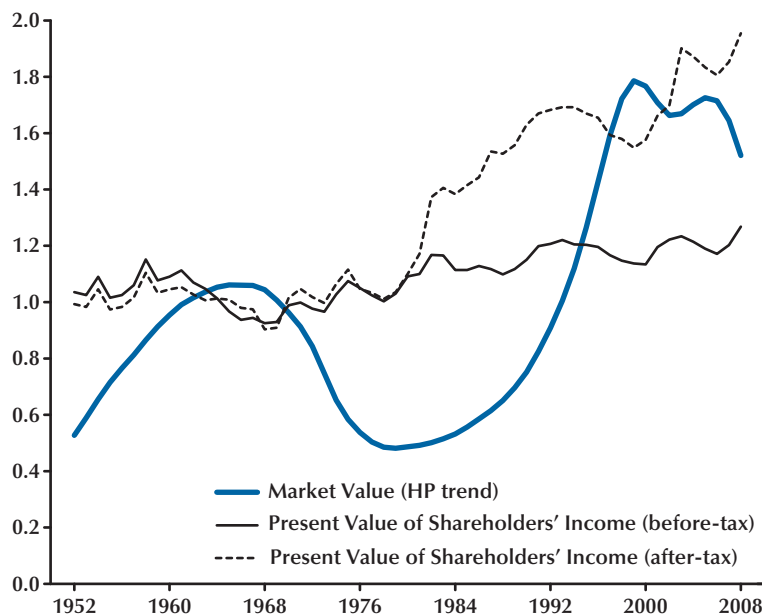
share) fluctuates widely over the sample period and has gone through the roof during the past decade. This fraction increased by 50 percent between 1953 and 1965, and then declined by 40 percent between 1966 and 1971 and remained at that level up to the early 1980s. For the period 1982-86, the data display a doubling of the capital income share, followed by a relative stabilization up to 2001, when the share increases sharply to unprecedented levels. By 2007 the capital income share is 40 percent higher than its 1983-2001 average and more than double what it was during the late 1960s and the 1970s. So much for the widespread assumption of long-run constancy of the capital share. Notice also that the long-run fluctuations we capture by means of the HP filter trend are dwarfed by the fluctuations taking place at business cycle frequencies: Factor shares in corporate income are anything but stable over time.

As a matter of fact, fluctuations in capital income are so large (in particular the 1982-2007 increase) that it becomes meaningless to perform a perfect foresight experiment symmetric to the

ones noted above using the historically realized average growth rate of the capital share as the out-of-sample predictor for future dividends' growth. Because of the very large growth in the share of value added that is devoted to capital, the average growth rate of corporate capital income during the past 25 years is close to 6 percent, while corporate GDP grew on average by 3 percent. Assuming a permanent growth rate of 6 percent for corporate capital income implies that its share of corporate value added would become 100 percent a few decades in the future, which clearly makes no sense. A more reasonable experiment would then be to assume that, in the future, the capital share would remain constant at its average level during, say, the past 10 years, while corporate GDP grows into the infinite future at some reasonable rate. This is what we do, assuming that the capital share will remain at its average of the period 1998-2008 and, for consistency with our previous analysis, assuming that the current dividend tax will persist into the infinite future. The results are summarized in Figure 9. Interesting enough,

Figure 9

Present Value of Shareholders' Income versus Market Value (Relative to Corporate GDP)



this reasonable modification does not make much of a difference, and the simulated market ratio resembles that of Figure 5.


It is also instructive to consider what the second equality implied by the theory, and by common sense, suggests: In normal circumstances the market value of a corporation should be equal to the replacement cost of its installed capital stock *plus* whatever organizational and intangible capital (patents, industrial secrets, and so on) the corporation controls. In principle, at least, a corporation should be able to sell its buildings, structures, and equipment (i.e., its tangible assets) on the market at roughly their replacement cost; hence, its market valuation should be lower than that of its tangible assets only in those special circumstances in which such tangible assets have been poorly invested and cannot be redirected to a different productive activity. Although this happens all the time at the level of the individual firm, one does not expect this to happen for roughly 32 of 56 years for the entire corporate

sector. This is exactly what these data suggest happened!

Specifically, the time-series evolution of the K/Y ratio (at replacement value) as reported by NIPA, and summarized in Figure 2, moves a lot less than the market ratio; and it seems to be strongly negatively correlated with it as well. We observe very high investment levels (hence, of the capital stock in relation to output) in the mid-1970s, while equity values are persistently low. When the stock market trend inverts, so do investment and the K/Y ratio—hence, the low investment in the 1980s, with high equity values. Put differently, until about 1987, whenever an American firm purchased a piece of capital and installed it, that piece of capital immediately lost some value according to the stock market's prices. The common sense interpretation of this fact is that, for more than 30 years, the stock market considered the investment decisions of U.S. corporations to be “value reducing”! We call this a “puzzle” and, unless one is willing to theorize that “negative organizational capital” was accu-

Figure 10

Trends in Market Value and in Shareholders' Income (Relative to Corporate GDP)



mulated for three decades, this puzzle dwarfs the many others.

A recent literature offers an alternative interpretation of the previous facts. A high level of investment takes place when new and profitable technologies are first discovered, or adopted due to changes in the economic environment; profits come in later, when those technologies become fully operational and start producing fruitful results. Moreover, the fact that new technologies and new capital are introduced may render old capital obsolete, causing the market value of the latter to collapse (compare with Hobijn and Jovanovic, 2001, and Peralta-Alva, 2007). In this sense, one is tempted to read the high profits after 1982 as the return on the high level of investment of the 1970s. While this interpretation is perfectly reasonable and makes historical sense, it does require us to throw away a major tenet of most standard models: that is, that the stock market's prices embed an unbiased forecast of future corporate performances. The depressed stock mar-

ket values of the mid-1970s and early 1980s did not seem to incorporate the payoffs of the ultimately successful investments of this period. Hence, these payoffs cannot be conceived as "expected." They happened, but the shareholders financing the high investment levels of those years were apparently unable to foresee the future gains those investments would have brought to them. This is puzzling.

During those years, instead, the share of capital in corporate income was at historical lows (Figure 10) and the market ratio seemed to, myopically, reflect more current miseries than future successes. The fact that it is hard to reconcile these observations with the theory is also emphasized by the technological change–driven explanations for the trends in the market ratio quoted above, because in those models the market ratio tends to recover much earlier than in the data. Notice that it is only after the mid-1980s, when firms' successes had been observed for some time and the capital share of corporate income has

started to rise steadily, that the market ratio also picks up and starts reflecting current successes—or, maybe, forecasting future ones. A similar point can be made for pretty much every single major swing of the data we are considering: *Oscillations in the market ratio are anticipated by oscillations in the share of capital income in corporate GDP.* This observation suggests the need to look more carefully into the way in which shareholders forecast future performances and, in particular, into the role that current movements in the share of capital income play in determining shareholders' optimism or pessimism for the future.

BUILDING ON THE LESSONS FROM STANDARD MODELS

The main conclusions we derive from our previous analysis are as follows. First, actual dividends paid are too smooth to account for the key low-frequency trends in the market ratio; this remains true also when unexpected changes in the fiscal regime are taken into consideration. However, actual dividends paid are not necessarily what is priced by the stock market: While dividends paid are stable over time, we documented that the fraction of corporate value added captured by shareholders displays important fluctuations. Despite this adjustment, the asset valuation equation implied by different models under the perfect foresight hypothesis and a constant interest rate predicts a market ratio that is still too smooth relative to the data. Furthermore, we find that large classes of asset-pricing models where dividends are endogenously determined have some predictions that are hard to reconcile with the data. In particular, these models predict that market value should equal the value of the assets of the firm (after adjusting for taxes); in the data, these two series are negatively correlated—and most of the time the market value of the firm is lower than the value of the physical assets the firm controls! Finally, eyeball analysis suggests that movements in the share of capital in corporate income may be rough but consistent predictors of movements in the market ratio, an empirical regularity we now try to exploit.

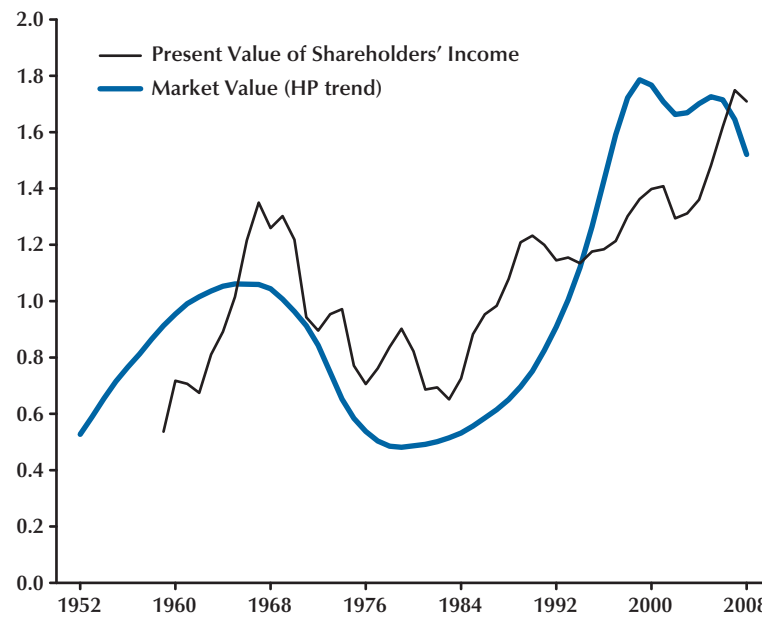
A delicate issue with all of our previous computations is the following: Pretending that in 1950 or in 1960, or even 1992 for that matter, shareholders could exactly forecast dividend payments in 2007 is clearly absurd. More important, the perfect foresight assumption typically results in a smooth series of predicted market values since all future fluctuations are foreseen and capitalized from the very beginning. Hence, a more reasonable “expectations formation” hypothesis needs to be introduced. While doing this opens a bottomless can of worms, this is an issue one must face squarely, especially if the study of past stock market behavior is supposed to shed some light on its current performance: What on earth drives shareholders' expectations? We consider this issue next.

In the past, people talked of “extrapolative expectations,” arguing that—when forecasting the future in the absence of an understanding of the structural model driving the system—we look at trends in past data and extrapolate those trends over the relevant horizon. This happens, though, only when we have become convinced that they are, indeed, permanent trends and not just small and irrelevant blips. When evidence suggests that the trend has changed or reverted to old patterns, we accept it only after a while; but, once accepted, we tend to extrapolate it into the indefinite future. The problem, obviously, is how long is the “while” and how reasonable is it to assume that people extrapolate trends that cannot be sustained forever, such as the one we just noted to exist in the capital share of corporate value added during the past two decades? This is a hard question for which we do not have a good answer and that the learning literature has really never addressed. We will try, nevertheless, to make some practical progress along these lines.

We start by assuming that people extrapolate past trends forward, altering them as soon as “enough” evidence is obtained that the previous trend is no longer likely to persist. Following this idea, suppose agents use all the information available up to T periods in the past. We generate separate “forecasts” for the growth rate of corporate GDP and for the capital income share, using a weighted average on the observations for the

Figure 11

Present Value of Shareholders' Income (Distributed Lags) versus Market Value (Relative to Corporate GDP)



last $N < T$ periods. We focus, as before, on the classical trading strategy where infinitely long series of shareholder income are generated and used to predict market value. We compute “forecasts” based on the distributed lag equation,

$$X_{t+1} = \frac{1-\lambda}{1-\lambda^{N+1}} \sum_{i=0}^N \lambda^i X_{t-i},$$

where X is the growth rate of each variable under consideration (in this case, the capital income share, corporate GDP growth rate, and dividend income tax rates).

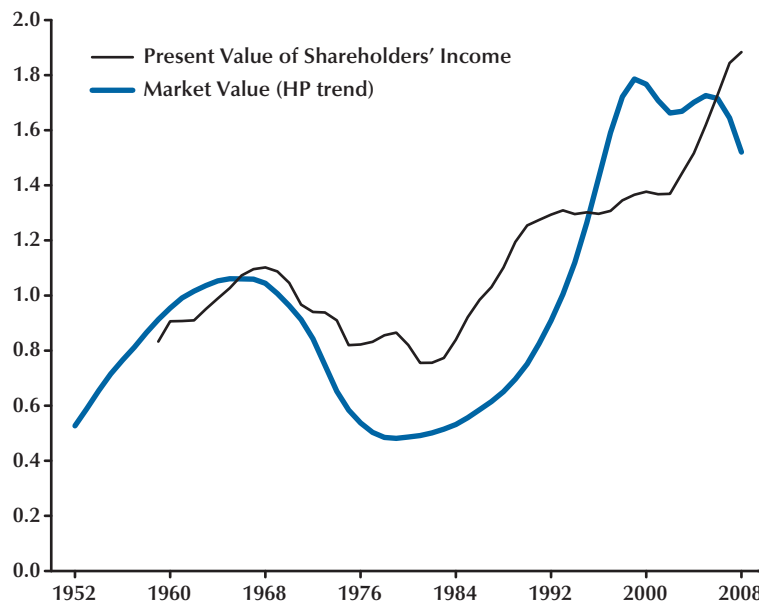
We assume, as in our previous quantitative experiments, a constant discount factor of 7 percent and use the maximum number of lags possible at each moment in time (given our dataset). We then estimate the weights (one lambda for each time series being forecasted) in the distributed lag equation such that the sum of squared deviations between the theoretical market ratio and the data is minimized. (Further details can be found in our computational appendix.) The model’s pre-

dictions and the data are summarized in Figure 11. As we can see, this simple approach can deliver a substantial improvement over the perfect foresight framework considered before. In particular, the predicted magnitudes for the 1960-68 increase, the mid-1970s’ decline and stagnation that followed, and the ultimate recovery of market valuations are all comparable to those in the data. Notice, however, that the timing of the predicted changes in the trend of the market ratio tend to be off by a few years and that we cannot account for the large drop in market value of recent years either. Nevertheless, given the simplicity of our approach, we consider the predictions obtained by using this ad hoc form of extrapolative expectations interesting and worthy of more systematic pursuit.

Barsky and De Long (1993) follow a similar approach and conclude that dividend movements roughly account for the secular fluctuations in the U.S. stock market from the 1800s to 1993. These authors, however, abstract from changes in taxes

Figure 12

Present Value of Shareholders' Income (Distributed Lags and Constant Long-Run Growth) versus Market Value (Relative to Corporate GDP)



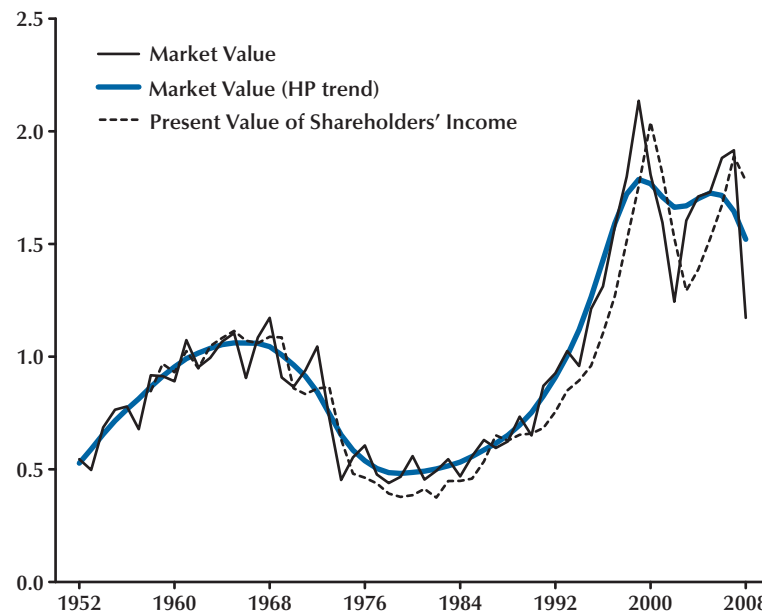
and use actual dividends paid by stock market firms in their analysis. As we illustrate above, however, actual dividends paid cannot account for the high stock market values observed after 1993. More important, these authors use a distributed lag equation similar to ours and estimate, period by period, a *permanent* growth rate of dividends. Such an estimation process implies (when applied to the actual data) that agents must expect dividends to grow at a rate permanently higher, or lower, than corporate GDP, which cannot really happen. Barsky and De Long’s paper uses data up to the very early 1990s and hence does not have to face this puzzling prediction of their methodology, which is instead an implication of the past two decades of data. Furthermore, our quantitative analysis above has documented that, once one constrains the long-run growth rate of dividends to equal that of corporate GDP, it becomes impossible to account for observed stock market fluctuations. Our results in Figure 11 are computed based on forecasts of the HP trend of the growth rate of corporate GDP and, although

this growth rate is far less volatile than dividends, may thus be subject to a similar criticism. To determine whether our results are driven by potentially unrealistic, permanent, forecasted values for the growth rate of corporate GDP, we consider a new experiment where a constant 3 percent growth rate for corporate GDP is assumed throughout. The results are summarized in Figure 12.

The fit of the model is still surprisingly good, but only up to a year ago: The drop in market ratio of the past year was unpredictable on the basis of the dividends performance observable up to 2007. It remains an open question as to what this methodology would predict a couple of years from now, when the substantial drop in capital income that took place in 2008 and 2009 will be reported in the data. This scenario notwithstanding, though, Figure 12 suggests that fluctuations in model-consistent shareholder income, and in taxes, can account for a large part of the secular movements of the U.S. stock market during the past 56 years *if* one is willing to assume that market participants use something akin to the distributed-lags fore-

Figure 13

Present Value of Shareholders' Income (Selling Shares After a Period) versus Market Value (Relative to Corporate GDP)



casting equation above in forming their expectations about the future. The key challenge for this simple framework seems to be accounting for the timing of the recovery during the mid-1980s and early 1990s. Capital income increased dramatically in the early 1980s. Dividend taxes declined substantially during the mid-1980s as well. According to the theory, these changes should have translated in a strong stock market recovery at the time. In the data, the recovery did start in the mid-1980s, but most of it did not take place until the mid-1990s.

Up to now, we have evaluated the asset-pricing equation of the basic model under a trading strategy of buy and hold (forever). Notice, however, that this fundamental asset-pricing equation holds for buy and hold, but it also implies that the value of the firm must equal the value of holding shares for *any* number of periods, T , and then selling and capturing the corresponding capital gains (or losses). Unfortunately, our current framework of analysis is not suited for studying the implications of the theory for short-term trading strategies.

To understand why, observe that such analysis would require forecasting model-consistent capital income (as before), as well as future market values. But model-consistent capital income is a relatively small fraction of corporate GDP (between 6 and 10 percent), while market value is almost 10 times larger, between 50 and 160 percent of corporate GDP. Hence, for relatively short holding periods, fluctuations in the value of corporations predicted by the theory will be mostly driven by fluctuations in predicted market values. Indeed, when we apply the previous methods to derive the predictions of the asset-pricing equation for holding periods between three and five years, we obtain a very good fit not only for the HP trend of the market ratio, but also for the actual market ratio (Figure 13). The fact that the model matches well the HP trend of the market ratio, however, follows immediately from the fact that the HP trend of the market value (which is the focus of our analysis) is very persistent and predictable. To compute the predictions of the theory for

period t 's market value, we assume agents use all information available up to when the forecast is made, compute a new HP trend for market value, and use this market value trend to forecast (using the distributed lag equation above) future market values (and thus capital gains). Since an HP trend series that is updated continuously provides a good approximation (with a lag) to the underlying time series, our estimation method approximates well the actual market ratio (with a lag).

CONCLUDING REMARKS

We study fluctuations in the long-run trend of the ratio between stock market value and GDP for the U.S. corporate sector. According to economic theory, the market value of U.S. corporations should equal the *expected* present discounted value of the *future* flow of income and capital gains generated by this sector. This prediction of the theory is frequently tested assuming *perfect foresight* on *actual* dividends paid. Actual dividends are very smooth and their movements cannot account for stock market trends, even in the long run. Many researchers consider this a puzzle. We find that a measure of *model-consistent* dividends fluctuates much more than actual dividends paid. More important, fluctuations in model-consistent dividends are positively correlated with stock market fluctuations. We illustrate that the perfect foresight assumption, by construction, predicts a very smooth present value of model-consistent dividends, and thus a very smooth market ratio, even when dividends fluctuate a lot. Theory does not require and does not imply that individuals and firms have perfect foresight, however; it simply requires and predicts that individuals will use all available information optimally (that is, as well as they can) to form their expectations of future movements in capital income. We then evaluate the theory under the assumption that all available (but no future) information is used in an extrapolative expectations format to forecast future dividend payments. We use a distributed lag equation to do so. We find that the present value of dividends, computed in this way,

is much more consistent with the data. Apart from the obvious question of what, other than wisdom after the fact, may justify or explain the particular choice of forecasting rule made by market participants, our analysis leaves open an important puzzle: The value of corporations should equal the value of their tangible and intangible assets, while in the data the two series seem to be negatively correlated and remain persistently apart from each other.

REFERENCES

- Barsky, Robert B. and De Long, James B. "Why Does the Stock Market Fluctuate?" *Quarterly Journal of Economics*, May 1993, 108(2), pp. 291-311.
- Bian, Jiangze. *Three Essays on Macroeconomics and Asset Prices*. Ph.D. Dissertation, University of Miami, 2007.
- Boldrin, Michele; Christiano, Lawrence J. and Fisher, Jonas D.M. "Habit Persistence, Asset Returns, and the Business Cycle." *American Economic Review*, March 2001, 91(1), pp. 149-66.
- Campbell, John Y. "Consumption-Based Asset Pricing," in G.M. Constantinides, M. Harris and R.M. Stulz, eds., *Handbook of the Economics of Finance*. New York: Elsevier, 2003, pp. 803-997.
- Easterbrook, Frank H. "Two Agency-Cost Explanation of Dividends." *American Economic Review*, September 1984, 74(4), pp. 650-59.
- Feldstein, Martin and Green, Jerry. "Why Do Companies Pay Dividends?" *American Economic Review*, 1983, 73(1), pp. 17-30.
- Hobijn, Bart and Jovanovic, Bojan. "The Information-Technology Revolution and the Stock Market: Evidence." *American Economic Review*, December 2001, 91(5), pp. 1203-20.
- McGrattan, Ellen R. and Prescott, Edward C. "Taxes, Regulation, and the Value of U.S. and U.K. Corporations." *Review of Economic Studies*, July 2005, 72(252), pp. 767-96.

Peralta-Alva, Adrian. “The Information Technology Revolution and the Puzzling Trends in Tobin’s Average ‘q.’” *International Economic Review*, August 2007, 48(3), pp. 929-51.

CEPR Discussion Paper No. 2858, Centre for Economic Policy Research, 2001.

Shiller, Robert J. *Irrational Exuberance*. Princeton, NJ: Princeton University Press, 2005.

Ravn, Morten O. and Uhlig, Harald. “On Adjusting the HP-Filter for the Frequency of Observations.”

DATA APPENDIX

- The *market value of corporations* is based on quarterly data in “Table L.213: Issues at Market Value” from the Flows of Funds Account of the United States from the Federal Reserve. In a given year, we take the last quarter of the quarterly market value to create the annual level.
- *Corporate value added (or corporate GDP)* is “Table 1.14: Gross Value Added of Corporate Business” from the National Income and Product Accounts (NIPA) published by the Bureau of Economic Analysis (BEA). Data on corporate businesses is obtained NIPA.
- The *replacement value of corporate capital* is the sum of nonresidential and residential tangible corporate fixed assets, measured at current cost, as reported in the Standard Fixed Asset Tables (Tables 4.1 and 5.1) published by the BEA.
- *Dividend tax rates* up to 1998 are from the data appendix of McGrattan and Prescott (2005), and the rate from 1999-2005 from Bian (2007), who followed the methodology of McGrattan and Prescott. The rate for 2006-2008 was assumed equal to that for 2005.
- Both *compensation of employees* and *corporate income taxes* were taken from “Table 1.14: Gross Value Added of Domestic Corporate Business” in Domestic Product and Income from NIPA. Compensation of employees is line 4, while taxes are the sum of corporate income taxes (line 12) and taxes on production and import less subsidies (line 7). Corporate *gross investment* is from the Standard Fixed Asset Tables (Tables 4.7 and 5.7), and it is equal to the sum of Investment in Private Nonresidential and Residential Fixed Assets of U.S. corporations.

Transformations on the Data

Except for tax rates, we divide all the series by the GDP implicit price deflator (2000 = 100) before we begin our analysis. In the annual series to which we apply the HP trend, we use a smoothing parameter of 6.25, recommended by Ravn and Uhlig (2001). To estimate Figure 13, we use the HP trend of the quarterly market value. In this case, we use a smoothing parameter of 1600.

COMPUTATIONAL APPENDIX

This section describes the algorithm used to compute the market value forecasts summarized in Figures 11 and 12.

We start by constructing forecasts for corporate GDP growth, the share of dividends in corporate GDP, and dividend tax rates. These forecasts are based on a distributed lag equation. In particular, standing at period t , we compute the period $t+1$ forecasted value for variable X as follows:

$$\tilde{X}_{t+1t} = \frac{1 - \lambda_x}{1 - \lambda_x^N} \sum_{i=0}^N \lambda_x^i X_{t-ilt}.$$

Here, λ_x denotes the weight of past observations, and N the number of lags included in the forecast. Let $(\tilde{g}_{t+k/t}, \tilde{\tau}_{d,t+k/t}, \tilde{s}_{d,t+k/t})$ denote the resulting period t -forecasted sequences for the growth rate of corporate GDP, the tax rate on dividend income taxes, and the dividend share of corporate GDP.

Then, our forecast at period t for dividends at $t+k$ is given by:

$$\tilde{d}_{t+k/t} = (1 - \tilde{\tau}_{d,t+k/t}) \tilde{s}_{d,t+k/t} GDP_{corp,HP-trend,t} \prod_k (1 + \tilde{g}_{t+k/t}).$$

We compute the predictions of the model using *only* information available up to the period when the forecast is made. This entails computing a new set of HP trends from the data, as well as new out-of-sample forecasted sequences for dividends, GDP growth, and taxes for each given year.

The market values reported in Figure 11 are thus given by the present value of dividends condition:

$$V_t = \sum_{k=1}^{\infty} \frac{\tilde{d}_{t+k/t}}{(1+r)^k}.$$

Finally, the values of the weights in the distributed lag equations above, $(\lambda_d, \lambda_{\tau_d})$, where chosen to minimize the square sum of residuals between forecasted and observed market values (as ratios to corporate GDP). The values we used are $(\lambda_d = 0.65, \lambda_{\tau_d} = 0.79)$.

The results reported in Figure 12 are obtained following a symmetric procedure, where we replace the last term of the dividend forecasting equation above,

$$\prod_k (1 + \tilde{g}_{t+k/t}),$$

by $(1 + 0.03)^k$.