possible. In summary, however, an important observation on the current state of the analyst literature is that it is almost exclusively based on indirect evidence.

The earliest research on financial analysts developed as a by-product of capital markets research focused on correlations between accounting earnings and stock prices. In that line of research, it was necessary to quantify the amount of 'news' in earnings announcements. Thus, a measure of 'expected' earnings was required, which was compared to earnings actually reported, allowing a quantification of the 'unexpected' component of earnings. In an informationally efficient market, this unexpected news should lead to immediate short-window stock price reactions.

The interest in tests of market efficiency and value relevance of accounting earnings prompted a significant amount of research on time-series modeling of earnings. This literature is extensive and generated much discussion about then new topics in the accounting literature such as earnings response coefficients (ERCs), ARIMA parameters, impulse response functions, and so on. This literature seems to have reached its peak during the late 1970s and early 1980s, at which time researchers gravitated towards using analysts' forecasts of earnings as a substitute for the complex time-series models. This launched a number of studies that ran horse races between analysts' forecasts and time-series models to see which was a better measure of the 'expected' component of earnings. Fried and Givoly (1982) are often given credit as the paper that supported the definitive conclusion that analysts are a better proxy for expected earnings than estimates from time-series models.

Although there remains scattered interest in the time-series properties of earnings, Kothari (2001) recently commented that the literature on time-series modeling of earnings is "fast becoming extinct … [due to] the easy availability of a better substitute: analysts' forecasts are available at a low cost in machine-readable form for a large fraction of publicly traded firms." As it became generally accepted that analysts' forecasts were superior to time-series forecasts, academics became interested in a deeper understanding of analysts' forecasts and analysts' themselves. Among academic accountants, analysts were elevated to the status of an economic agent in the capital markets worthy of extensive study. As a result, more recent work attempts to understand analysts' incentives, conflicts of interest, loss functions, and so on. Prior to briefly reviewing what we know about analysts, it is important to articulate why we *still* study analysts.

The cynical response to why academics still study analysts is that the data are easy and cheap to access. Several companies like First Call, I/B/E/S, Value Line, and Zacks maintain databases on the forecasts and recommendations of thousands of analysts covering thousands of companies, allowing easy use of these data by academic researchers. Perhaps an even more cynical response is that academics very much enjoy analyzing distributions (i.e., means, medians, standard deviations, etc.) and correlations. Analyst data are easily converted into variables that provide interesting distributions and correlations (e.g., signed forecast error, forecast accuracy, ERCs, etc.).

However, the real reason I believe research on analysts continues is that we are interested in how the capital markets function, and examining analysts furthers such knowledge. On one hand, analysts are one of the preeminent market information intermediaries, distributing forecasts and results of their analysis to institutional and individual investors. Thus, examining properties of the analysts' forecasts and analysis helps us understand the nature of the information that seems to be impounded in stock prices. Another perspective is that analysts are a good proxy for beliefs held by investors in general, so examining properties of analyst data provides insight into how investors in general utilize and process accounting information like financial statements, footnotes, and other financial disclosures. Finally, having elevated analysts to the status of an interesting set of economics agents for detailed study, it is intrinsically interesting to study what analysts do and how they utilize financial accounting information. This final reason explains most of the current work on analysts.

## OVERVIEW OF WHAT WE KNOW (OR THINK WE KNOW)

Early survey research and anecdotal evidence suggest that analysts are voracious for all kinds of information (e.g., Tevelow 1971, Chandra 1974, Frishkoff, Frishkoff, and Bouwman 1984, Epstein and Palepu 1999). It is not surprising, however, that in responding to surveys, analysts would tend indicate they always prefer more information to less. It is one thing to simply express a desire for information and another to incur costs to acquire or process it, particularly given a drastic increase in the length of annual reports in recent years (Li 2006). Research on analysts' information needs and preferences is generally regarded as 'descriptive' and is frequently overlooked in empirical research. This is unfortunate, because investigations on what information analysts might use and how they use it should incorporate these findings, if for no other reason than to see if what analysts say is consistent with what it appears they actually do.

Prior to discussing specific observations on generally accepted findings in the literature, a very brief discussion of the evolution of the literature is in order. Figure 2 provides a timeline that highlights general trends in the literature between the 1960s and early 2000s. Let me again emphasize that this is not meant to be a literature review or a comprehensive summary of all primary questions examined. Additionally, figure 2 is employed as a heuristic to place the subsequent discussion of specific observations in context. The reader is directed to the literature reviews identified in the introduction for a full list of questions and a more comprehensive coverage of relevant studies. Also, I will provide very brief highlights of each paper, and the brevity of these oversimplified highlights will necessarily oversimplify and undersell the full contribution of the paper.

As previously discussed, the initial impetus for examining analysts forecasts was the need for a better proxy for earnings expectations to be used in capital markets research. This literature spanned approximately two decades (1968-1987) and appears in the lower left quadrant of figure 2. Brief highlights of notable conclusion from these studies are as follows:

- Cragg and Malkiel (1968): Five-year growth rates forecasted by analysts were no different than simple algebraic extrapolations.
- Elton and Gruber (1972): Annual forecasts by various groups (pension fund, investment advisors, investment bank analysts) were no different between naïve time-series model and each group of analysts.
- Barefield and Comiskey (1975): Analysts' forecasts outperformed a simple no-change earnings forecast model.
- Brown and Rozeff (1978): Analysts' forecasts outperformed 'less naïve' time-series models, especially at longer forecast horizons.
- Fried and Givoly (1982): Using a (then) large sample of panel data (100 forecasts per year for 1969-1979), analysts' forecasts were more accurate than those from various time-series models.
- Brown, Griffin, Hagerman, and Zmijewski (1987): Analysts' forecast superiority over time-series models is due to (i) a timing advantage and (ii) an information advantage.

These studies primarily appeared in finance journals, employed small samples relative to those typical in current analyst research (e.g., hundreds of observations vs. hundreds of thousands), and used research designs that ran horse races between different forecasts. Fried and Givoly (1982) is generally recognized as having provided the most compelling evidence that analysts are superior to time-series models and several years later, Brown et al. (1987) clarified the source of analysts' superiority. Thus, it took almost two decades for researchers to settle comfortably on the conclusion that analysts were better than time-series models at forecasting earnings. However, as discussed below, the economic magnitude of analysts' superiority appears to be small, suggesting that analysts' value to the capital markets likely rests on other roles than simply forecasting earnings.

Building on the research that compared analysts relative to time-series models, research considered refinements and extensions to research designs, with the goal of identifying factors that are correlated with incremental earnings forecast accuracy. These studies also appear in the lower left quadrant of figure 2, and are briefly highlighted below:

- O'Brien (1988): The most recent forecast more accurate than consensus.
- O'Brien (1990): There is no evidence of an analyst-level effect on forecast accuracy, thus no analysts are persistently better than others.
- Stickel (1990): Analysts ranked as an Institutional Investor All-Star are superior forecasters than a matched sample based on forecast recency.
- Brown (1991): The accuracy of the consensus forecast gets more accurate if older forecasts are dropped.
- Sinha, Brown, and Das (1997): Careful controls for forecast recency yield evidence that some analysts are more accurate than others
- Mikhail, Walther, and Willis (1997): Individual analyst experience increases forecast accuracy
- Clement (1999): Analysts' forecast accuracy is increasing in resources and decreasing in complexity.

Thus, the literature moved beyond concern over analysts being superior to time-series models, and began investigating whether some analysts were better than others. As with the previous efforts on analysts versus time-series models, this series of research initially showed no differences, but subsequently found the existence of differences.

Simultaneous to these two sets of studies, research was also considering the association of analysts' forecasting activities with stock prices. Some of the papers highlighted above also examined market reactions to forecasts and earnings surprises. For example,

- Fried and Givoly (1982) and others: Earnings forecast accuracy generally corresponds to a greater association between unexpected earnings based on such forecasts and announcement period stock returns.
- O'Brien (1988): Even though Standard & Poors and I/B/E/S analysts exhibit higher forecast accuracy, they have no stronger association with stock returns than time series models.
- Philbrick and Ricks (1991): The actual definition of what income statement level earnings being forecasted varies across forecast data providers. Value Line forecast errors are the smallest, but various combinations of forecasts and actual earnings across the databases yields the strongest association with announcement period stock returns (e.g., unexpected earnings based on Value Line earnings forecasts and I/B/E/S actual earnings)

This focus on the correlation between analysts-based earnings surprises and stock prices prompted researchers to examine whether analysts' themselves appeared to be efficient with respect to information cues. Such studies tend to examine whether analyst forecast errors are correlated with publicly available information. If a correlation exists, research concludes that analysts are inefficient with respect to such information. This area of research arose around 1990 and continues to the present. Studies shown in the top right quadrant of figure 2 are highlighted below:

- De Bondt and Thaler (1990): Analysts overreact to past earnings changes, resulting in forecasts that are overoptimistic.
- Lys and Sohn (1990) and Abarbanell (1991): Analysts' forecasts underreact to information in prior stock price changes.
- Mendenhall (1991) and Abarbanell and Bernard (1992): Analysts underestimate the serial correlation in quarterly earnings (i.e., post-earnings announcement drift), but to a lesser extent than investors do through stock prices.
- Elliott, Philbrick, and Wiedman (1995): Analysts systematically underreact to their own sequential prior forecast revisions.
- Easterwood and Nutt (1999): Analysts underreact to negative information and overreact to positive information, both reactions leading to analysts being persistently overoptimistic.
- Bradshaw, Richardson, and Sloan (2001): Analysts underreact to predictable earnings patterns following extreme accruals.

As can be seen from the highlights, there does not appear to be a general consensus on whether analysts over- or underreact to information. Either way, the conclusions that are inevitably that analysts are 'inefficient' with respect to numerous pieces of information. This literature is vast, with almost any information cue one can consider having been subjected to an analyst forecast analysis. In the next section, I argue that drawing conclusions about the efficiency of analysts' forecasts based on correlations may not be a strong test of analysts' processing of information.

A second wave of research on the efficiency of analysts attempts to understand whether analysts are internally efficient with respect to their own information outputs. For example, given the correspondence between earnings expectations and value, do analysts efficiently use their own earnings forecasts in valuing companies and generating stock recommendations? Select papers include:

- Bradshaw (2004): Analysts' recommendations are consistent with the use of heuristic valuations incorporating their own earnings forecasts.

- Asquith, Mikhail, and Au (2005): Qualitative information in analysts' reports explains a significant amount of their recommendations, target prices, and the price reaction to these forecasts.
- Loh and Mian (2006): More accurate forecasts lead to more profitable stock recommendations.

This research is noteworthy in that it necessarily considers simultaneously more outputs from the analyst than just the earnings forecasts. As argued in the next section, the literature on analysts suffers from an overemphasis on earnings forecasts relative to other important tasks performed by analysts. In this spirit, many of what some consider to be the most interesting papers on analysts focus on their activities within the context of what their individual and employer-level incentives are. A sampling of these types of papers is as follows:

- Francis and Philbrick (1993): Analysts trade off earnings forecast accuracy for intentional optimism to curry favor with managers.
- McNichols and O'Brien (1997): Analysts' exhibit a self-selection bias such that negative views are censored, and hence unobservable to investors or researchers.
- Lin and McNichols (1998): Analysts exhibit overoptimism when their employers perform investment banking services for covered firms.
- Michaely and Womack (1999): After the quiet period following an initial public offering, affiliated analysts are more likely to issue buy recommendations than are unaffiliated analysts.
- Mikhail, Walther, and Willis (1999): Forecast accuracy is negatively related to analyst job turnover.
- Hong and Kubik (2003): Promotions and demotions at investment banks depend more on optimism than accuracy.
- Gu and Wu (2003) and Basu and Markov (2004): These papers question analysts' loss functions implied by prior work that uses ordinary least squares models to link forecast errors and various measures (implying a quadratic loss function) by proposing that analysts' might prefer to minimize the absolute error instead.
- Raedy, Shane, and Yang (2006): Evidence of analyst underreaction might not be due to them ignoring publicly available information, but due to their asymmetric loss function whereby they incur greater reputation cost

of forecast errors when the error has the opposite sign as the analysts' prior earnings forecast revision. (i.e., bad to 'overshoot').

Left out of the terse listing of papers in figure 2 are many important studies on (i) the analyst coverage decision, (ii) dispersion and its association with prices and accuracy, (iii) recent changes in the regulatory environment (FD), and (iv) experimental research that has a bearing on decision processes (but I'll defer discussion of these until later). I have also focused the studies listed here on those involving earnings forecasts, which is consistent with the representativeness of earnings forecasts as the focus of most studies in this literature. It is only recently that researchers have begun investigating recommendations (Womack 1996), growth projections (LaPorta 1996), and target prices (Brav and Lehavy 2003).

The overall takeaways from the above discussion is that approximately four decades of research on analysts focuses heavily on the earnings forecasting task, with only recently increasing interest in other activities performed by analysts. Second, the literature moves relatively carefully, with the conclusion that analysts dominate time-series models taking two decades. Third, beginning in the 1990s, much work has been positioned as attempts to understand what information analysts use and how they use it (i.e., the black box). Finally, as research studies have begun to consider activities beyond basic earnings forecasting, it has become necessary (and interesting) to examine analysts' incentives and investigate what role they might play in the empirical regularities developed over the past several decades of research (e.g., optimism). The next section provides ten specific observations that may guide future thought on how to interpret and advance the evidence on analysts' and their roles in the capital markets.

## SPECIFIC OBSERVATIONS ON WHAT WE KNOW (OR THINK WE KNOW)

### 1. Analysts' Forecasts are Optimistic

Of all the regularities regarding sell-side analysts, the understanding that analysts' forecasts are routinely optimistic is the most pervasive. Numerous studies document that analysts' forecasts of earnings end up, on average, being too high. The problem is that this is a sweeping generalization that is not *on average* descriptive. There are at least three qualifications to the generalization that analysts are routinely optimistic. First, what specific forecasts are believed to be optimistic – quarterly earnings per share forecasts, annual earnings per share forecasts, growth forecasts, target prices, sales forecasts, cash forecasts, etc.? The typical explanation for why analysts would be persistently optimistic is that they wish to maintain cordial relationships with management, and optimistic forecasts further this goal. However, with regards to the most prevalent forecast made by analysts, earnings per share, it is difficult to understand why the managers analysts are presumably trying to please would prefer optimistic earnings forecasts. Research makes it clear that forecast errors (measured as actual earnings minus the forecast) are positively correlated with stock price reactions. Thus, forecasts that are too high (i.e., optimistic) create negative forecast errors and negative stock price reactions. On average, managers would seem to desire avoiding such reactions. Indeed, recent evidence in the accounting literature examines the 'meet or beat' phenomenon, which describes the preference by managers and tendency for quarterly earnings announcements to equal or slightly exceed analysts' forecasts. Overall, it appears that at least for short-term forecasts, it is not descriptive to generalize that analysts' forecasts are optimistic.

Second, we seem to be well aware of selection biases in analyst forecast data which form the basis of most of our research. Several studies indicate that analysts seem to follow the old adage, 'if you don't have anything good to say, don't say anything at all.' For example, analysts are reluctant to issue negative recommendations (i.e., 'sell'), and more important, having issued favorable recommendations, they exhibit a reluctance or sluggishness in downgrading recommendations. Even though this is a well-known phenomenon, we apparently disregard knowledge of this selection bias in drawing generalities about the overall level of analyst optimism. In other words, what is interpreted as persistent optimistic bias by analysts could simply reflect the fact that we do not get to observe analysts' pessimistic views. With the recent implementation of NASD 2711 and NYSE 472 rules that, among other things, require analyst research reports to provide benchmark distributions of the brokerage's recommendations and target prices, we may witness an increasing tendency for analysts to convey previously non-communicated pessimistic views.

Finally, a recent body of research on 'street' or 'pro forma' earnings has revealed issues with analyst forecast data that systematically result in optimistically biased forecasts. Firm managers have always highlighted earnings in earnings releases that exclude the effect of various one-time charges. However, this practice escalated beginning in the 1990s, and firms began reporting earnings excluding an even greater number of income statement line items, including, for example, research and development expense, advertising expense, customer acquisition costs, and so on. As these examples suggest, the types of income statement amounts excluded were disproportionately expenses (rather than gains or revenues). Both Bradshaw and Sloan (2002) and Abarbanell and Lehavy (2007) note that forecast data providers such as First Call and I/B/E/S claim to archive actual earnings figures that match the earnings definition being forecasted by the majority of analysts. This is important because the standard practice to calculate analyst forecast error (and hence bias) is to subtract the actual earnings figure from the forecast database from the forecast. Thus, if analysts forecast earnings before the effects of one-time items and research and development expense, then the forecast data providers include the actual earnings before one-time items and research and development expense in the historical database used by academics. Evidence presented in both papers referenced above indicate that the forecast data providers seem to have only gradually adjusted the actual earnings figures on the database to correspond to figures being forecasted by analysts. Both papers identify 1992 as representing a marked shift in the correspondence of actual and forecasted earnings. As much of the research supporting the inference that analysts are persistently optimistic was published using pre-1992 data, the non-correspondence between the actual earnings used in those studies (i.e., bottom-line 'net income' from Compustat or one of the forecast data providers) would have systematically resulted in mechanically upwardly biased forecast errors.

## 2. Analysts' Forecasts Are Superior to Time-Series Model Forecasts

The second presumably well-known feature of analysts' forecasts is that they are superior to forecasts from time-series models. Accounting research aimed at modeling earnings using ARIMA models was at its peak during the 1970's and seems to have effectively ended in the mid-1980's. Brown (1993) provides a comprehensive review of much of this literature, which is also briefly summarized by Kothari (2001), who states at the outset (p. 145), "I deliberately keep my remarks on the earnings' time-series properties short because I believe this literature is fast becoming extinct. ... [due to] easy availability of a better substitute: analysts' forecasts...."

On one hand, if analysts are efficient in any sense, as has been noted before by Brown et al. (1987), it has to be the case that analysts' forecasts outperform time-series model forecasts, because analysts have both a timing and information advantage. Analysts can easily calculate any anointed time-series model and incorporate that information into their overall information set. Moreover, because time-series models are parsimonious, the information available to analysts is greater than that which can be quantified by any time-series model. Thus, for most forecast dates, an analyst will have an information advantage over a time-series model, which necessarily relies on historical inputs. Nevertheless, it took scores of papers spanning two decades (i.e., approximately 1968-1987) for academic research to conclude that analysts' are superior to time-series models.

Many of the papers that concluded examined the relative forecasting ability of analysts versus time-series models were based on limited samples. For example, Barefield and Comiskey (1975) examine forecasts for 100 firms (and conclude that analysts outperformed a simple random walk forecast) and Brown and Rozeff (1978) examine forecasts for 50 firms (and conclude that most time-series models are outperformed by analysts, particularly at longer horizons). Fried and Givoly (1982) is generally credited as one of the decisive studies in this area, primarily due to the significantly expanded sample size. They examine 100 forecasts per year for the period 1969-1979 and conclude that analysts were superior to time-series models. However, what seems to have been overshadowed in subsequent research that wholly abandoned time-series models is the slim margin by which analysts won this contest. For example, Fried and Givoly calculate absolute forecast errors scaled by actual earnings per share. Their primary results indicate an average absolute forecast error for analysts of 16% relative to a comparable forecast error for two time-series models of 19% and 20%, respectively. Furthermore, results for individual years are often closer than this 3-4% spread. This seems to be a slim margin of victory for analysts given the information and timing advantages they have over the time-series models. The increasing tendency for managers to provide earnings guidance (Matsumoto 2002) and earnings preannouncements (Soffer, Thiagarajan, and Walther 2000) should have increased analysts' superiority over time-series models, but no research of which I am aware has examined this.

If one restricts their consumption of research to accounting journals, then it would appear that research using time-series models is indeed extinct.[1] However, outside of the accounting literature, continued use of time-series forecasts as an alternative and as a benchmark for expert forecasts is prevalent. Indeed, the economics literature largely concludes that time-series forecasts are superior to those of various experts. For example, this is argued to be the case for forecasts of interest rates (Belongia 1987), gross domestic product (Loungani 2000), recessions (Fintzen and Stekler 1999), and business

[1] This is not meant to dispute the conclusion in Kothari (2001) referenced above, which is indeed accurate.

cycles (Zarnowitz 1991). This discrepancy in conclusions across research paradigms is surely related to the unit of analysis. Forecasts of earnings is done frequently with the input of the preparers of the earnings being forecasted, accounting procedures for those earnings are well-understood, and such accounting standards often have the objective of smoothing reported earnings (e.g., pension assumptions). In contrast, items like interest rates, GDP, recessions, and business cycles are not generally subject to the control of an individual manager or follow a prescribed set of rule governing their reporting.

### 3. Analysts' Forecasts are Inefficient

A large number of research papers spanning the late 1980s through the present examine whether analysts' forecasts are 'efficient.' Similar to how efficient market prices are defined, forecasts are said to be efficient if they incorporate all information available to the analyst. Thus, studies have examined whether analysts incorporate information in past earnings, past market prices, and past forecast revisions; similarly, more recent studies examine whether analysts' forecasts are efficient with respect to information in financial statement information like accruals, management forecasts, and various other financial disclosures.

These studies inevitably draw conclusions about the efficiency of analysts' forecasts. If forecast errors are correlated with some information available *ex ante* to the analyst, the forecast is said to be inefficient with respect to that information. In these cases, the analyst is said to have either 'underreacted' or 'overreacted' to the information. As it turns out, it is rare to witness empirical results which support an efficient use of information. The likely reason is that the data we rely upon is noisy, which inevitably leads to coefficients in empirical tests that are consistent with inefficient use of information.

To clarify this, consider a simple correlation between some analyst variable AV (e.g., annual forecast revision) and some variable of interest X (e.g., information in a quarterly earnings announcement). What the researcher wants to measure is corr(AV, X). However, X is likely measured with error, so the researcher ends up measuring X+error, rather than X. In the typical regression framework, the researcher would estimate the following regression:

$$AV = \alpha + \beta(X+\text{error})+e,$$

leading to the well-known downward bias in the estimate of $\beta$ (absent other covariates). This downward bias inevitably leads researchers to conclude that, with respect to the information in the phenomenon measured by X, analysts appear to be inefficient. The often overlooked or unstated alternative is that the tyranny of measurement error contaminates our ability to draw strong conclusions regarding analysts' efficiency in processing particular pieces of information.[2]

## 4. Most Academic Research Ignores Analysts' Multi-Tasking

Of the hundreds of papers published on sell-side analysts, casual empiricism supports the conclusion that most focus exclusively on the earnings forecasting process. Thus, if someone unfamiliar with sell-side analysts went to the accounting and finance

[2] Of course, if the left hand side were some analyst variable, like forecast error, measurement error would tend to bias this simple univariate specification towards a conclusion of efficiency rather than inefficiency. The variety of empirical specifications in the literature and the multivariate (rather than simple univariate) nature of such specifications leads to ambiguous directional predictions regarding measurement error induced bias, but it is reasonable to presume that conclusions that generally fall between full efficient use of information by analysts and complete inefficiency are most likely.

literature to understand what it is they do, they would likely come away with the impression that analysts' primary goal is to issue accurate earnings per share forecasts.

In contrast, consideration of all the roles performed by an analyst suggests that earnings per share forecasts are either tangential or at best just one of many inputs into the analysts' other (primary) activities. Thus, a focus on earnings forecasts by academics is useful to understanding what analysts do, but it is a means not an end. Schipper (1991) noted early on in this literature that, "The general focus of accounting research on accuracy and bias of analysts' earnings forecasts has yet to capitalize on whatever opportunities for insights might arise from considering these forecasts in the context of *what the analyst does* … [emphasis added] (p. 112). Similarly, Zmijewski (1993) argued shortly thereafter that one of the primary areas of research that could further our knowledge are studies that lead to "expansion of our analysis of financial analysts' earnings forecasts to encompass more of what they actually do [emphasis added] (p. 338).

The easiest means of understanding what analysts do is to examine other outputs provided by them. In recent years, research into these other outputs has been growing, with studies on stock recommendations (e.g., Womack 1996), growth projections (e.g., Dechow and Sloan 1997), target prices (e.g., Brav and Lehavy 2003), and risk ratings (Lui, Markov, and Tamayo 2007). A second step is to simultaneously examine these outputs. In other words, if one of analysts' primary objectives is to issue an investment recommendation for a security, then one might examine how earnings forecasts and growth projections are associated with the actual recommendation (e.g., Bradshaw 2004). To gather a quick feel for how active research is along these suggestions, I performed a global search of scholarly articles on ABI/INFORM using various keywords, and found the following:

| | |
|---|---|
| analyst+earnings | 867 articles |
| analyst+recommendation | 149 articles |
| analyst+long+term+growth | 54 articles |
| analyst+target+price | 14 articles |
| analyst+earnings+recommendation | 27 articles |
| analyst+earnings+long+term+growth | 22 articles |
| analyst+earnings+target+price | 3 articles |
| analyst+earnings+recommendation+long+term+growth | 1 article |

This is not to suggest that research studies that incorporate more than one analyst variable are superior, but rather, that furthering our understanding of what analysts do and why they do it requires consideration of their portfolio of activities. For example, Loh and Mian (2006) examine whether analysts who provide superior earnings forecasts also provide more profitable stock recommendations, which is a useful question to answer as it pertains directly to the use of earnings forecasts as an input into the arguably more important role of providing investment advice.

Clearly, as discussed above, the overwhelming bulk of research effort appears to focus on earnings forecasts, with some distant level of interest on analysts' stock recommendations. However, beyond that the interest level suggested by the above ABI/INFORM search seems to drop substantially. The simple explanation may simply be that data on these other metrics have not been widely available until recently. For example, whereas large samples of machine-readable earnings forecast data have been available since the early 1970s, data for long-term growth forecasts became available in 1981, for recommendations in 1992, and for target prices in 1996. I return to this theme later when I comment on research that is aimed at understanding what analysts' do with their own earnings forecasts.

## 5. Analysts are Dominated by Conflicts of Interest

Besides the first point raised regarding the belief that analysts' forecasts are persistently overoptimistic, perhaps the second most prevalent belief is that analysts' behavior is dominated by conflicts of interest. There are at least six sources of conflicts that have been discussed either in the literature or the financial press and that are purported to lead to analysts being overoptimistic. The following briefly lists, in my assessment, the sources of conflict in descending order of the relative emphasis given to them in the literature.

1. <u>Investment banking fees.</u> Managers periodically require access to the capital markets and require the assistance of investment banking professionals, who are frequently employed by firms that also run sell-side research shops. It has long been argued, and recent anecdotal evidence is consistent with the charge, that sell-side research departments are rewarded by the investment banking side of operations for providing favorable coverage of deals that the firm underwrites. Such fees are the fuel of such firms, and typical large placements bring in millions of dollars in fees. Accordingly, sell-side research, which is generally a cost rather than a profit center, is argued to be predisposed towards overoptimism due to the lure of lucrative investment banking fees. This explanation is the most prevalent.

2. <u>Currying favor with management</u>. Distinct from the incentive to appease managers to obtain investment banking business, sell-side analysts have also been accused of being optimistic so that they maintain access to firm managers who are a primary source of information flow (Francis and Philbrick 1993). The recently implemented Regulation FD is meant to curb this practice, and requires that managers refrain from selectively releasing private information. Several studies have attempted to examine whether the implementation of this regulation led to less optimistic forecasts and recommendations by analysts. However, around the same time that Regulation FD was implemented, there were other regulations and market sentiment changes that make it difficult to attribute any observed change in overall analyst optimism to this single piece of regulation (e.g., NYSE 472, Nasdaq 2711, Sarbanes-Oxley, large interest rate changes, severe currency

exchange changes, etc.). Even in the presence of regulation disallowing selective disclosure, there remain reasons for analysts to maintain cordial relations with managers (e.g., simply getting managers to return phone calls, receiving favorable queuing during conference calls, etc.).

3. Trade generation incentives. Another reason analysts are allegedly predisposed towards optimism is that their firms also receive compensation through handling investor trades. As the argument goes, it is easier to convince an investor to buy a stock that they do not own rather than convincing them to sell a stock they must already own. Consequently, to generate investor purchases, analysts will optimistically bias their reports. Recent evidence by Cowen et al. (2006) and Jacob et al. (2008) suggests that incentives for optimistic bias are stronger for trading than for investment banking. They partition investment banks into those that provide investment banking and those that do not, where trading fees are the primary source of revenues, and find that *ex post* optimistic bias is stronger for analysts working at the non-investment bank firms. Also, Jacob et al. (2008) provide some evidence that affiliated analysts are actually more accurate than unaffiliated analysts, and moreover, the differential forecast accuracy appears due to the employment of better analysts and the presence of greater resources.

4. Institutional investor relationships. The close ties between institutional investors and investment banks also provide sources of conflicts for sell-side analysts. As recipients of sell-side research, institutions may take positions in securities based on the information and recommendations conveyed in analysts' formal reports. If an analyst then downgraded a security that an institution had taken a position in, this would clearly be viewed unfavorably by the institution.

5. Research for hire. Given that approximately one-third of public companies have no analyst coverage and over half have at most two analysts, a recent phenomenon in equity research is for companies to pay for research to be conducted on their company. Several consortiums have been established, such as the National Research Exchange and the Independent Research Network. The conflicts of interest in these arrangements are obvious, and it remains to be seen how these will be managed.

6. Themselves. Finally, an often overlooked source of conflicts for analysts is the behavioral bias inherent in the analysis of securities. Similar to the well-documented home bias in the finance literature, the familiarity analysts develop with firms and their managers can lead analysts to develop close affinity to a firm.

This affinity may then result in analysts seeing the firm ‘through rose-colored glasses,’ and being incapable of downgrading or forecasting negative outcomes.

Of these six sources of analyst conflicts, the allegation that lucrative investment banking fees is the most cogent. Clearly, regardless of the reputation of a particular investment bank, any right-minded manager would steer clear of their services if sell-side analysts employed by that investment bank held negative views on the firm. Researchers have investigated such effects extensively, and it would appear that most researchers subscribe to the belief that these conflicts have strong effects on observed optimism in analysts’ reports. Numerous studies document significantly more optimistic forecasts and recommendations for affiliated analysts (e.g., Lin and McNichols 1998, Michaely and Womack 1999, Dechow, Hutton, and Sloan 2000, Lin, McNichols, and O’Brien (2005).

One explanation other than analysts’ deliberate optimism inspired by investment banking business is that among the distribution of investment banks, some will be the employers of analysts that are more optimistic about a particular firm, and it is the selection of those investment banks by the managers that explains the documented optimism by affiliated analysts. Research is unable to distinguish between these two explanations, but Ljungqvist, Marston and Wilhelm (2006) offer some evidence consistent with management choice. They examine investment banking deal flows and find no evidence that overoptimistic recommendations by analysts explain investment banking selection, the main determinant being the strength of prior investment banking relationships. Another explanation is that there is a collective level of heightened positive sentiment about firms that are in the growth stage and hence need external financing. Consistent with this, Bradshaw, Richardson, and Sloan (2006) document that both affiliated and unaffiliated analysts display increasing optimism around periods of external financing and both groups show declines in the levels of optimism subsequent to external financing. This is not inconsistent with investment banking conflicts leading to optimism in research, but it does attenuate the degree of sinister interpretation given to the reports of analysts that are viewed as 'affiliated.' If analysts (as well as other market participants) tend to be optimistic about subsets of firms, it is not surprising that it would be the subset that is growing and seeking external financing.

However, it is instructive to review the economic significance of investment banking conflicts as documented in the literature. Lin and McNichols (1998) provide one of the most compelling studies to review because of the relatively large sample and well-executed matched sample design. They examine approximately 2,400 seasoned equity offerings (SEO) spanning 1989-1994. Primary results examine for significant differences in one-year ahead and two-year ahead earnings per share forecasts, growth projections, and stock recommendations. A summary of their results is as follows:

| | One-year ahead EPS | Two-year ahead EPS | Earnings growth | Stock Recommendation |
|---|---|---|---|---|
| Unaffiliated | 0.071 | 0.098 | 0.207 | 3.901 |
| Affiliated | 0.070 | 0.099 | 0.213 | 4.259 |
| *Difference* | -0.001 | 0.001 | 0.006 | 0.358 |
| Significant difference? | No | No | Yes | Yes |

*Note: EPS forecasts are scaled by price. Earnings growth projections reflect forecasts of annual percentage growth. Stock recommendations are coded on a 1 to 5 scale, with 1 being 'strong sell' and 5 being 'strong buy'.*

They find no differences in optimism in earnings forecasts, but they find analysts affiliated with SEOs provide higher growth projections and more positive recommendations. However, the economic significance of the differences do not seem large. For annual earnings growth projections, the difference is less than one percent, and the difference in stock recommendations is approximately one-third of a change in ranking. Adherents to the paradigm arguing that investment banking biases analysts to be optimistic would highlight that the analysts that are unaffiliated are almost as optimistic as the affiliated analysts because they too were using research to court the managers for the investment banking business, which is in conflict to the evidence discussed earlier in papers like Jacob et al. (2006).

**6. Limited Evidence Exists Regarding What Analysts Do with Their Own Forecasts**

It is presumed that analysts are sophisticated and their analyses are internally consistent. However, very little research has examined their outputs in a multivariate setting. For example, research has examined analysts' forecasting abilities extensively, and there have been moderate efforts to understand their recommendation abilities. Clearly, recommendations should be linked in some manner to analysts' valuations, and we believe from many capital markets studies (i.e., Ball and Brown 1968, etc.) that earnings expectations are positively correlated with prices. Thus, rational behavior by analysts would mean that their own earnings forecasts are correlated with their valuations that provide the basis for their stock recommendations.

Francis and Philbrick (1993) provided the earliest systematic study of the interplay between analysts' various forecasts. Although their sample prevents an examination of how individual analysts use their own forecasts. Nevertheless, their study is one of the first to attempt to understand how analysts incorporate specific information into their forecasts. They examined Value Line analysts, who issue earnings forecasts but include in their reports a 'timeliness ranking' of a stock, akin to an individual analyst's stock recommendation but prepared by other analysts at Value Line. They hypothesized that analysts would attempt to curry favor with managers by diffusing unfavorable timeliness rankings by optimistic forecasts, and they conclude that Value Line analysts appear to behave in this manner.

Another early study that attempted to directly examine the within-analyst correlation of various outputs is Bandyopadhyay, Brown, and Richardson (1995), who examine analysts' target prices and earnings forecasts. Based on the presumption that analysts use their own forecasts in deriving stock valuations, they hypothesize that both one-year ahead and two-year ahead earnings forecasts will be correlated with analysts target prices (i.e., valuations), and that the correlations will be stronger for longer horizon forecasts. Indeed, they document $R^2$s of approximately 30% (60%) when correlating changes in target prices with changes in one-year ahead (two-year ahead) earnings forecasts. Similarly, Loh and Mian (2006) find that analysts with more accurate earnings forecasts provide more profitable stock recommendations, consistent with analysts using their own forecasts as inputs into their valuations and recommendations.

Recently, there seems to be a growing understanding of the benefits of understanding analysts' use of information, and attempts to measure within-analyst correlations of data are becoming more common. For example, Bradshaw (2002) performed a content analysis and found that analysts' valuations are almost always based on various earnings-multiple heuristics, and Bradshaw (2004) documented that researcher-generated recommendations based on simple residual income valuations using analysts' earnings forecasts as inputs outperform the analysts' recommendations that are based on heuristics. Similarly, Barker (1999) and Asquith, Mikhail, and Au (2005) document a high degree of reliance by analysts on qualitative factors in communicating their analyses, supplementing their heuristic use of earnings forecasts to assess valuations of firms. Given increasing availability of line item forecasts other than earnings, there is also an increasing interest in the internal consistency of those measures as well. For example, Ertimur, Mayew, and Stubben (2008) examine the multiple-level forecast accuracy of analysts that provide disaggregated forecasts (i.e., sales and earnings).

The trend towards research that simultaneously considers multiple analyst outputs is a step in the right direction if our goal is to increase our knowledge of analysts using large sample databases. One of the common objectives of research on analysts is to provide evidence that allows us to peer inside the decision-making processes they follow. However, though there are benefits from the typical archival empirical approach, the methodology is necessarily limited in its ability to garner insights into how analysts make decisions. Alternatively, research methodologies that work with data other than the databases provided by I/B/E/S and other providers are likely to provide complementary approaches. The next two sections expand on these

## 7. We Think We Know *How* Analysts Forecast

As the literature on analysts has grown, researchers have moved beyond straightforward investigations of distributional properties of forecast errors and profitability of analysts' recommendations. The tenor of most studies is that the researchers are interested in *how* analysts perform their tasks. However, with few exceptions, none provide direct evidence on *how* analysts go about generating forecasts or making stock recommendations. The problem appears to be a preference for archival research, which is subject to data and methodological constraints. Thus, researchers tend towards similar approaches and typically regress forecast errors on different independent variables to explain forecast errors. Some papers attempt to provide indirect evidence, but the nature of these analyses limits the strength of conclusions we can draw about analysts' actual decision processes.

The typical research design adopted when a researcher holds some hypothesis about how analysts use some information signal is to estimate a regression of analyst forecast error on the information variable,

$$\text{Forecast Error} = \alpha + \beta X + e,$$

where X is the variable of interest. As summarized in figure XX, right-hand side variables have included past earnings changes, past price changes, analysts' forecast errors, income statement line items, balance sheet line items, financial statement footnote information, management forecasts, macroeconomic variables, and so on. From these econometric analyses, conclusions are drawn as to whether the analyst incorporated the information captured by the variable X in their earnings forecast process.

Such a research design is a study of associations, not behavior. However, it has become prevalent to draw conclusions regarding analysts' behavior from these tests. Notwithstanding the fact that the combination of the research designs and the conclusions do not actually speak to analysts' behavior, these results do not map into the way that forecasting is covered in most financial statement analysis courses and textbooks. This suggests that either the research designs that are utilized in an attempt to see into the forecasting process or the pedagogical approach to prospective analysis needs revision. At a minimum, it is important for researchers to be careful about drawing strong conclusions about analysts' behavior based only on data that can be quantified and used as inputs in a specification like that above.

One alternative is to continue the trend in simultaneously examining multiple analyst forecasts and other information, as discussed earlier. Though limited by the research design that relies on archival data, this approach allows extended insights into statistical associations. Combined with prior findings of associations between forecast errors and various information signals, multivariate analyses of analysts' outputs can address numerous interesting questions (e.g., does forecasting cash flows lead to more accurate forecasts, more profitable recommendations, and so on). The second alternative is to embrace alternative research methodologies, discussed next.

## 8. Empiricists Have Traditionally Not Embraced Alternative Methodologies (but This is Changing)

As noted above, the primary methodology employed in the analyst literature is the empirical analysis of archival data. With a few exceptions, only recently have other methodologies received more attention in the literature. A likely explanation for the disproportionate focus on analysis of archival data is that it is much less costly to download a panel of I/B/E/S data than it is to conduct an experiment or perform a content analysis of a distribution of analyst reports. This explanation mirrors the likely explanation for the disproportionate analysis of earnings forecast data relative to other analyst outputs for which data availability is lower, such as risk ratings and target prices.

An early paper by Larcker and Lessig (1983) is a good example of the limitation of statistical analysis of archival data. In this study, Larcker and Lessig perform an experiment with 31 subjects who were asked to make buy or no-buy decisions for 45 stocks. They were interested in the competing ability of linear modeling (i.e., regression analysis) and retroactive process tracing (i.e., ex post interviews of subjects) to accomplish two objectives: (i) predicting subjects buy and no-buy decision and (ii) identifying the relative importance of various information cues used by the subjects. These objectives continue to map very well into those of many analyst studies that employ archival data.

They found that both linear models and process tracing performed reasonably well at predicting the buy and no-buy decisions of the subjects. However, there were frequent differences between the two approaches in identifying relative cue importance to the subject's buy and no-buy decisions. These findings lead the authors to conclude that if the goal of a research study is the *prediction of a judgment decision*, then both approaches appear valid, and lower cost and complexity would favor linear modeling. However, if the goal of a research study is *to understand what information is used and how it is used*, a technique like retroactive process tracing seems necessary. This point cannot be emphasized enough, as it bears directly on the 'black box' in figure 1b.

The current shortcoming of the literature on sell-side analysts is our lack of understanding of what goes on inside the black box of what an analyst actually does. Fortunately, there is a growing use of alternative methodologies that complement research that uses linear models. Alternative approaches to understanding analysts' activities include surveys and interviews, experiments, rigorous content analysis approaches, and focused analysis of representative firms). Clearly, alternatives to linear modeling also have weaknesses (i.e., surveys risk biased responses, experiments have difficulty replicating complex unstructured tasks, content analysis only has access to the final communication medium rather than the process itself, analyzing a single brokerage firm may have no external validity, etc.). For such reasons, these approaches are to be viewed as complementary. Together, consistent evidence across alternative methodologies increases validity of research conclusions and is necessary for this literature to progress.

The popularity of the recent survey of managers by Graham, Harvey, and Rajgopal (2005) is testament to the level of potential interest in the results of a survey of financial executives. Although there are a number of various surveys of financial analysts, most are relatively limited in scope or geography.[3] A notable exception is a survey by Block (1999), who surveyed members of the Association for Investment Management and Research (AIMR). His survey was broadly focused and queried analysts on their uses of valuation models, importance of financial inputs, bases for recommendations, various opinions regarding market efficiency and dynamics. The most remarkable finding in his survey is that analysts overwhelmingly do not emphasize present value models to value firms. Additionally, he found that analysts do not pay much attention to dividend policy, they focus more on the long-term prospects than near-term quarterly results, and analysts believe that skilled portfolio managers can beat the market.

[3] For example, surveys have focused on analysts' opinions of cash flow accounting (McEnroe 1996) and forecast revisions (Moyes, Saadouni, Simon, and Williams 2001), and have been conducted in various international markets including Saudi Arabia (Alrazeen 1999), Japan (Mande and Ortman 2002), Belgium (Orens and Lybaert 2007), and China (Hu, Lin, and Li 2008).

As noted above, surveys provide useful insights, but a weakness is the possibility that respondents do not truthfully report. However, as also noted above, if this survey evidence is combined with alternative research methodologies and the results consistently point towards the same conclusion, concerns over threats to validity can be minimized. As an example of how a conclusion can be compelling based on the collective results from studies using alternative methodologies, consider the conclusion in Block (1999) that analysts do not rely very much on present value models. This could be due to some form of non-response bias, a miscommunication of what was meant by present value techniques, or analysts' concerns that their approaches are proprietary and they bias their responses. However, subsequent studies that adopted content analysis (Bradshaw 2002) and linear modeling (Bradshaw 2004) provide uniformly consistent results that analysts indeed do not appear to make stock recommendations consistent with present value-based models.

Published surveys on analysts are relatively rare, as are content analyses and focused studies of individual brokerage firms. Moreover, those that are published appear to be concentrated outside of what are typically considered 'top-tier' journals. This is unfortunate, because other than my own personal interactions with analysts and users of analysts' information, where most of my knowledge of analysts has been obtained, I have learned a great deal from reading these studies. On an optimistic note, research utilizing experimental research methods is much more common and seems to be increasingly acceptable to top-tier journals. Many of these types of studies employ undergraduate or graduate students as subjects, but it is becoming increasingly common to see actual analysts serving as subjects. For example, Libby et al. (2008) employ a sample of 81 experience analysts and examine the tension between maintenance of relationships with firm managers and optimism and pessimism in earnings forecasts. Perhaps more interesting than the actual experimental results, the post-experiment subject interviews provide insights into how analysts are aware of the optimism-to-pessimism pattern in earnings across fiscal periods, but believe this pattern helps them receive preferential treatment in conference calls. Again, echoing the theme that multiple research designs can be combined to increase the validity of a research conclusion, the evidence in Libby et al. (2008) regarding analysts' desire to receive preferential or favorable treatment in conference calls (even in a post-Regulation FD environment) is also shown by Mayew (2008), who extracted data from conference call transcripts. His archival empirical study also confirms that analysts' with optimistic research on a company get more attention during conference calls. Together the Mayew and Libby et al. studies give increased comfort that analysts are indeed still concerned about currying favor with managers.

A final trend that is serving to make research on analysts more cohesive across methodologies is a growing prevalence of accounting academics properly trained in experimental research techniques. Moreover, this is accompanied by the gaining acceptance of 'behavioral finance' research, which is incorporating psychology research on decision making. The majority of experimental accounting research relies on similar theories (Koonce and Mercer 2005). Further, researchers appear to be realizing that certain methodologies are suited for specific research questions. For questions which arise around situations of decision-making and information processing, experiments seem useful because of their ability to minimize confounding 'real-world' variables and manipulate the variables of interest (Bloomfield, Libby, and Nelson 2002).

## 9. Academics May Be Focusing Too Much on the Least Important Activities

As has been noted, the vast majority of research on analysts is focused on their ability to forecast earnings. The early literature pitted analysts against time-series forecasts, then gravitated towards identifying superior analysts with more accurate earnings forecasts. Recently, researchers have been simultaneously considering the interplay among various analyst outputs (e.g., earnings and recommendations), but the anchor of the analysis remains earnings forecast accuracy. If an individual with no understanding of sell-side analysts were to attempt to understand what they do based on a reading of our academic literature, that person would surely conclude that one of the things most important to analysts is their earnings forecasts. I contend that this would be a gross mischaracterization of the analyst's job function, and hence his/her incentives. I believe such a view characterizes that of many academics, and as a result impedes our ability to further our understanding of sell-side analysts.

To provide some perspective on the importance of earnings forecasts, table 1 provides a panel of data reflecting traits of analysts ranked in order of importance by respondents to the annual Institutional Investor Ranking of analysts. This ranking is the first-order determinant of an analyst's compensation (Groysberg, Healy, and Maber 2008). Thus, if we assume that analysts wish to maximize their compensation, then providing institutional investors with what they need, as reflected in the rankings, will be descriptive of aspects of their job towards which they devote significant effort.

The data in table 1 span 1998-2005, and show that the number of criteria reported in the rankings each year range from a low of eight items in 1998 to fifteen during 2002-

2004. The rankings indicate that the most important trait valued by institutional investors is industry knowledge, which has been the number one trait for all years of the survey. Clearly, analysts' are valued for their ability to see individual companies within the context of the industry as a whole. Other traits appear relatively stable in their importance across recent years, with two exceptions – earnings forecast and stock selection. Whereas earnings forecasts were ranked fifth in importance in 1998, they are ranked last in the most recent year in table 1. Similarly, stock selection was ranked as high as second in 1998, but has fallen to second-to-last in the last year of table 1. As a statistical measure of whether these changes are meaningful, table 2 provides a simple test of whether the changes in the ranking are significant. The mean change in rank is calculated for the annual changes in ranking, where rankings are converted to a [0,1] interval.[4] For both earnings forecast and stock selection traits, the average change in ranking across 1998-2005 is significantly negative, indicating that both measures have become less important to institutional investors, and presumably less important to analysts, relative to other characteristics. Of course, one explanation is that earnings forecasts and stock selection are viewed as necessary by institutional investors, and presumably by analysts as well, but that other aspects of their jobs are relatively more important. This is consistent with earnings forecasts and stock selection being important; however, as suggested above, it also is consistent with these aspects of an analyst's job being relatively unimportant when their roles are viewed in context.

---

[4] Each ranking is converted to RANK' to span the interval [0,1] as

$$RANK' = ( (NRANK+1)-RANK)/NRANK,$$

where NRANK is the number of characteristics listed in the annual ranking and RANK is the numerical rank of the characteristic. Characteristics ranked in other years but not on the ranking in any individual year are assigned RANK'=0.

I believe that part of our focus on earnings forecast accuracy is driven simply by the wide availability of data on analysts' earnings forecasts and actual earnings and a predilection of accounting academics towards the investigation of phenomena that can be quantified. Measuring the accuracy of an earnings per share forecast suits our comfort zone. Similarly, measuring recommendation profitability is also appealing, despite numerous alternative measurement criteria decisions (i.e., return accumulation period, raw or adjusted returns, etc.). What is a lot more difficult to measure is the measurement of important aspects of the analysts' job function such as industry knowledge, assessment of firm strategy or quality of management, accessibility, the tone of their contextual reports, and so on. Nevertheless, researchers in this area must be open to alternative methodologies and data if the literature on analysts is to proceed in a meaningful way.

**10. Analyst Data are Indirectly Helpful to Other Work Examining the Functioning of Capital Markets**

In contrast to other critical points raised above, the following point is a commendation of research on analysts. As noted above, research on analysts has become pervasive with the elevation of analysts to a status of interesting economic agent worthy of individual examination. Comments numbered one through nine focus on this aspect of analysts. There is another very useful role of research using analyst data, which is that these data can provide insights into questions that arise in other capital market studies. Specifically, the identification and examination of asset pricing anomalies is an active area of research in the finance and accounting literatures. In the typical study, researchers demonstrate that future stock returns are systematically associated with information available *ex ante* (e.g., past earnings changes, past price changes, accounting accruals, insider trading, etc.). Such studies are always subject to the 'bad model' criticism, which argues that the correlation reflects an incomplete control for priced risk rather than a true asset pricing anomaly that can be costlessly arbitraged away.

Because of the difficulty of convincingly capturing priced risk (or priced risk factors), an alternative to addressing the bad model criticism is to use a research design that skirts the risk issue. Whereas capital market anomalies all pertain to how investors incorporate information into prices, and analysts' roles include the incorporation of information into their research, it is frequently useful to examine documented anomalies in the context of analysts' research. For example, as an extension of the seminal studies by Bernard and Thomas (1989, 1990) on the post-earnings announcement drift anomaly, Abarbanell and Bernard (1992) examine whether analysts incorporate the autocorrelation structure documented in the Bernard and Thomas papers into their forecasts. They find that similar to market prices, analysts underreact to prior earnings changes. Accordingly, critics that dismissed the post-earnings announcement drift anomaly as a mis-measurement of risk must also explain why the phenomenon shows up in a non-asset pricing setting. Similar analyses have been conducted with respect to the glamour anomaly (Frankel and Lee 1998), the January effect (Ackert and Athanassakos 2000), and the accruals anomaly (Bradshaw, Richardson, and Sloan 2001; Barth and Hutton 2004),

## CONCLUSION

In summary, we have learned a lot about analysts and their role in capital markets. However, research has focused on a narrow set of analyst outputs to draw conclusions regarding what analysts do and how they do it. Further, this research is largely limited to variables that can be quantified, there is limited but growing investigation of the co-determination of analysts' outputs, and there is a disproportionately large emphasis on what is likely a relatively unimportant activity – forecasting earnings. For this literature to progress, research that provides any kind of penetration of the 'black box' of how analysts actually process information should be encouraged, even if methods or approaches are imperfect.

This literature finds itself at an interesting juncture of time, with numerous recent shocks to the capital markets (e.g., Regulation FD, $1.4 billion SEC/state regulator settlement against ten large investment banks, a new independent brokerage research requirement, disclosure requirements of NASD Rule 2711 and NYSE Rule 472, and a trend towards paying for analyst coverage). Thus, there are numerous opportunities for the literature to progress if researchers move beyond the current prevailing paradigm of performing univariate analyses of earnings forecasts. Zmijewski (1993) discussed a literature review by Brown (1993), and echoed similar sentiments to those offered here. In commenting on the state of the literature at that time, he stated, "That is not to say, however, that researching the 'same old' issues using the 'same old' methodologies will be informative.... It will, naturally, become more and more challenging to identify interesting questions and to design interesting and meaningful empirical tests."

## REFERENCES


Abarbanell, J. 1991. Do analysts' earnings forecasts incorporate information in prior stock price changes? *Journal of Accounting and Economics* 14 (2): 147-165.

Abarbanell, J., and B. Bernard. 1992. Tests of analysts' overreaction/underreaction to earnings information as an explanation for anomalous stock price behavior. *The Journal of Finance* 47 (3): 1181-1207.

Abarbanell, J., and R. Lehavy. 2007. Letting the 'Tail Wag the Dog": The debate over GAAP versus street earnings revisited. *Contemporary Accounting Research* 24 (3): 675-723.

Abdolmohammadi, M., R. Simnett, J. Thibodeau, and A. Wright. 2006. Sell-Side analysts' reports and the current external reporting model. *Accounting Horizons* 20 (4): 375-389.

Ackert, L., and G. Athanassakos. 2000. Institutional investors, analyst following and the January anomaly. *Journal of Business Finance and Accounting* 27 (3-4): 469-485.

Alrazeen, M.H., 1999. “The Quality of Annual Financial Reports of Saudi Corporations: Users' Perceptions and Companies' Practices.” Ph.D. Thesis, Cardiff University.

Asquith, P., M. Mikhail, and A.Au. 2005. Information content of equity analyst reports. *Journal of Financial Economics* 75 (2): 245-282.

Ball, R., and P. Brown. 1968. An empirical evaluation of accounting income numbers. *Journal of Accounting Research* 6 (2): 159-178.

Bandyopadhyay, S., L. Brown, and G. Richardson. 1995. Analysts' use of earnings forecasts in predicting stock returns: Forecast horizon effects. *International Journal of Forecasting* 11 (3): 429-445.

Barefield, R., and E. Comiskey. 1975. The accuracy of analysts’ forecasts of earnings per share. *Journal of Business Research* 3: 241-251.

Barker, R.. 1999. Survey and market-based evidence of industry-dependence in analysts' preferences between the dividend yield and price-earnings ratio valuation models.” *Journal of Business Finance and Accounting* 26 (3-4): 393-418.

Barth, M., and A. Hutton. 2004. Analyst earnings forecast revisions and the pricing of accruals. *Review of Accounting Studies* 9: 59-96.

Basu, S., and S. Markov. 2004. Loss function assumptions in rational expectations tests on financial analysts' earnings forecasts. *Journal of Accounting and Economics* 38: 171-203.

Belongia, M. 1987. Predicting interest rates: A comparison of professional and market-based forecasts. *Federal Reserve Bank of St. Louis Review* March: 9-15.

Bernard, V., and J. Thomas. 1989. Post-earnings-announcement drift: delayed price response or risk premium? *Journal of Accounting Research* 27 (3): 1-36.

Bernard, V., and J. Thomas. 1990. Evidence that stock prices do not fully reflect the implications of current earnings for future earnings. *Journal of Accounting and Economics* 13: 305-340.

Block, S. B. 1999. A study of financial analysts: Practice and theory. *Financial Analysts Journal* 55 (4): 86–95.

Bloomfield, R., R. Libby, and M. Nelson. 2002. Experimental research in financial accounting. *Accounting, Organizations and Society* 27 (8): 775-810.

Bradshaw, M., S. Richardson, and R. Sloan. 2001. Do analysts and auditors use information in accruals? *Journal of Accounting Research* 39 (1): 45-74.

Bradshaw, M. 2002. The use of target prices to justify sell-side analysts' stock recommendations. *Accounting Horizons* 16 (1): 27-41.

Bradshaw, M., and R. Sloan. 2002. GAAP versus the street: An empirical assessment of two alternative definitions of earnings." *Journal of Accounting Research* 40 (1): 41-66.

Bradshaw, M. 2004. How do analysts use their earnings forecasts in generating stock recommendations?" *The Accounting Review* 79 (1): 25-50.

Bradshaw, M., S. Richardson, and R. Sloan. 2006. The relation between corporate financing activities, analysts' forecasts and stock returns. *Journal of Accounting & Economics* 42 (1-2): 53-85.

Brav, A., and R. Lehavy. 2003. An empirical analysis of analysts' target prices: Short-term informativeness and long-term dynamics. *The Journal of Finance* 58 (5): 1933-1968.

Brown, L., and M. Rozeff. 1978. The superiority of analyst forecasts as measures of expectations: Evidence from earnings. *The Journal of Finance* 33 (1): 1-16.

Brown, L., P. Griffin, R. Hagerman, and M. Zmijewski. 1987. Security analyst superiority relative to univariate time-series models in forecasting quarterly earnings. *Journal of Accounting and Economics* 9: 61-87.

Brown, L. 1991. Forecast selection when all forecasts are not equally recent. *International Journal of Forecasting* 7 (3): 349-356.

Brown, L. 1993. Earnings forecasting research: Its implications for capital markets research." *International Journal of Forecasting* 9: 295-320.

Brown, L., ed. 2000. *I/B/E/S Research Bibliography: Sixth Edition.* I/B/E/S International Inc., New York, NY.

Brown, P. 1993. Comments on 'Earnings forecasting research: Its implications for capital markets research'. *International Journal of Forecasting* 9: 331-335.

Chandra, G. 1974. A study of the consensus on disclosure among public accountants and securities analysts. *The Accounting Review* 49 (4): 733-742.

Clement, M. 1999. Analyst forecast accuracy: Do ability, resources and portfolio complexity matter? *Journal of Accounting and Economics* 27 (3): 285-303.

Cowen, A., B. Groysberg, and P. Healy. 2006. Which types of analyst firms are more optimistic? *Journal of Accounting and Economics* 41 (1-2): 119-146.

Cragg, J., and B. Malkiel. 1968. The consensus and accuracy of some predictions of the growth of corporate earnings. The *Journal of Finance* 23 (1): 67-84.

De Bondt, W., and R. Thaler. 1990. Do security analysts overreact? *The American Economic Review* 80 (2): 52-57.

Dechow, P., and R. Sloan. 1997. Returns to contrarian investment strategies: Tests of naïve expectations hypotheses. *Journal of Financial Economics* 43 (1): 3-27.

DeChow, P., A. Hutton, and R. Sloan. 2000. The relation between analysts'forecasts of long-term earnings growth and stock price performance following equity offerings. *Contemporary Accounting Research* 17 (1): 1-32.

Dunn, K., and S. Nathan, 2005. "Analyst Industry Diversification and Earnings Forecast Accuracy." *Journal of Investing* 14, 7-14.

Easterwood, J., and S. Nutt. 1999. Inefficiency in analysts' earnings forecasts: Systematic misreaction or systematic optimism? *The Journal of Finance* 54 (5): 1777-1797.

Elliot, J., D. Philbrick, and C. Wiedman. 1995. Evidence from archival data on the relation between security analysts' forecast errors and prior forecast revisions. *Contemporary Accounting Research* 11 (2): 919-938.

Elton, E., and M. Gruber. 1972. Earnings estimates and the accuracy of expectational data. *Management Science* 18 (8): 409-424.

Epstein, M., and K. Palepu. 1999. What financial analysts want. *Strategic Finance* 80 (10): 48-52.

Ertimur, Y., W.J. Mayew, and S.R. Stubben. 2008. Analyst reputation and the issuance of disaggregated earnings forecasts to I/B/E/S. Working paper, Duke University and University of North Carolina.

Fintzen, D., and H.O. Stekler. 1999. Why did forecasters fail to predict the 1990 recession? *International Journal of Forecasting* 15: 309-323.

Francis, J., and D. Philbrick. 1993. Analysts' decisions as products of a multi-task environment. *Journal of Accounting Research* 31 (2): 216-230.

Frankel, R., and C. Lee. 1998. Accounting valuation, market expectation, and cross-sectional stock returns. *Journal of Accounting and Economics* 25 (3): 283-319.

Fried, D., and D. Givoly. 1982. Financial analysts' forecasts of earnings: a better surrogate for market expectations. *Journal of Accounting and Economics* 4 (2): 85-107.

Frishkoff, P., P. Frishkoff, and M. Bouwman. 1984. Use of accounting data in screening by financial analysts. *Journal of Accounting, Auditing, and Finance* 8 (1): 44-53.

Givoly, D., and J. Lakonishok. 1984. Properties of analysts' forecasts of earnings: A review and analysis of the research. *Journal of Accounting Literature* 3: 117-152.

Graham, J., C. Harvey, and S. Rajgopal. 2005. The economic implications of corporate financial reporting. *Journal of Accounting and Economics* 40 (1-3): 3-73.

Groysberg, B., P.Healy, and D. Maber. 2008. What drives financial analyst compensation? Working Paper, Harvard Business School.

Gu, Z., and J. Wu. 2003. Earnings skewness and analyst forecast bias. *Journal of Accounting and Economics* 35 (1): 5-29.

Hirst, D., and P. Hopkins. 1998. Comprehensive income reporting and analysts' valuation judgments. *Journal of Accounting Research* 36: 47-75.

Hong, H., and J. Kubik. 2003. Analyzing the analysts: Career concerns and biased earnings forecasts. *The Journal of Finance* 58 (1): 313-351.

Hu, Y., T.W. Lin, and S. Li. 2008. An examination of factors affecting Chinese financial analysts' information comprehension, analyzing ability, and job quality. *Review of Quantitative Finance and Accounting* 30 (4): 397-417.

Jacob, J., S. Rock, and D.P. Weber, 2008. Do non-investment bank analysts make better earnings forecasts? *Journal of Accounting, Auditing and Finance*, Winter, Vol. 23 Iss. 1, 23-61.

Koonce, L., and M. Mercer. 2005. Using psychology theories in archival financial accounting research. *Journal of Accounting Literature* 24: 175-214.

Kothari, S.P. 2001. Capital markets research in accounting. *Journal of Accounting and Economics* 31: 105-231.

La Porta, R. 1996. Expectations and the cross-section of stock returns. *The Journal of Finance* 51 (5): 1715-1742.

Larcker, D.F., and V.P. Lessig. 1983. An examination of the linear and retrospective process tracing approaches to judgment modeling. *The Accounting Review* 58 (1): 58-77.

Li, F. 2006. Annual report readability, current earnings, and earnings persistence. Working paper, University of Michigan.

Libby, R., J. Hunton, H. Tan, and N. Seybert. 2008. Relationship incentives and the optimistic/pessimistic pattern in analysts' forecasts." *Journal of Accounting Research* 46 (1): 173-198.

Lin, H., and M. McNichols. 1998. Underwriting relationships, analysts' earnings forecasts and investment recommendations. *Journal of Accounting and Economics* 25: 101-127.

Lin, H., M. McNichols, and P. O'Brien. 2005. Analyst impartiality and investment banking relationships. *Journal of Accounting Research* 43 (4): 623-650.

Ljungqvist, A., F. Marston, and W.J. Wilhelm Jr. 2006. Competing for securities underwriting mandates: Banking relationships and analyst recommendations. *Journal of Finance*, Vol. 61, Iss. 1, 301-340.

Loh, R., and G. Mian. 2006. Do accurate earnings forecasts facilitate superior investment recommendations? *Journal of Financial Economics* 80 (2): 455-483.

Loungani, P. 2000. How accurate are private sector forecasts? Cross-country evidence from consensus forecasts of output growth. IMF Working Paper No. 00/77, April.

Lui, D., S. Markov, and A. Tamayo. 2007. What makes a stock risky? Evidence from sell-side analysts' risk ratings. *Journal of Accounting Research* 45 (3):629-665.

Lys, T., and S. Sohn. 1990. The association between revisions of financial analysts' earnings forecasts and security-price changes. *Journal of Accounting and Economics* 13 (4): 341-363.

Mande, V., and R. Ortman. 2002. Are recent segment disclosures of Japanese firms useful? Views of Japanese financial analysts. *The International Journal of Accounting* 37 (1): 27-46.

Matsumoto, D. 2002. Management's incentives to avoid negative earnings surprises. *The Accounting Review* 77 (3): 483-514.

Mayew, W. 2008. Evidence of management discrimination among analysts during earnings conference calls. *Journal of Accounting Research*, Forthcoming,

McEnroe, J.E. 1996. An examination of attitudes involving cash flow accounting: Implications for the content of cash flow statements. *The International Journal of Accounting* 31 (2): 161-174.

McNichols, M., and P.O'Brien. 1997. Self-selection and analyst coverage. *Journal of Accounting Research* 35 (Supplement): 167-199.

Mendenhall, R. 1991. Evidence on the possible underweighting of earnings-related information. *Journal of Accounting Research* 29 (1): 170-179.

Michaely, R., and K. Womack. 1999. Conflict of interest and the credibility of underwriter analyst recommendations. *The Review of Financial Studies* 12 (4): 653-686.

Mikhail, M., B.Walther, and R. Willis. 1997. Do security analysts improve their performance with experience? *Journal of Accounting Research* 35: 131-157.

Mikhail, M., B. Walther, and R. Willis. 1999. Does forecast accuracy matter to security analysts? *The Accounting Review* 74: 185-200.

Moyes, G., B. Saadouni, J. Simon, and P. Williams. 2001. A comparison of factors affecting UK and US analyst forecast revisions. *The International Journal of Accounting* 36 (1): 47-63.

Nichols, C., and J. Whalen. 2004. How do earnings numbers relate to stock returns? A review of classic accounting research with updated evidence. *Accounting Horizons* 18 (4): 263-286.

O'Brien, P. 1988. Analysts' forecasts as earnings expectations. *Journal of Accounting and Economics* 10: 53-83.

O'Brien, P. 1990. Forecast accuracy of individual analysts in nine industries. *Journal of Accounting Research* 28 (2): 286-304.

O'Hanlon, J. 1993. Commentary on: Lawrence D. Brown 'Earnings forecasting research: Its implications for capital markets research. *International Journal of Forecasting* 9: 321-323.

Orens, R., and N. Lybaert. 2007. Does the financial analysts' usage of non-financial information influence the analysts' forecast accuracy? Some evidence from the Belgian sell-side financial analyst. *The International Journal of Accounting* 42 (3): 237-271.

Philbrick, D., and W. Ricks. 1991. Using Value Line and IBES analyst forecasts in accounting research. *Journal of Accounting Research* 29 (2): 397-417.

Raedy, J., P. Shane, and S. Yang. 2006. Horizon-dependent underreaction in financial analysts' earnings forecasts. *Contemporary Accounting Research* 23 (1): 291-322.

Ramnath, S., S. Rock, and P. Shane. 2008. The Financial analyst forecasting literature: A taxonomy with suggestions for further research. *International Journal of Forecasting*, 24 (1): 34-75.

Schipper, K. 1991. Commentary on analysts' forecasts. *Accounting Horizons* 5 (4): 105-121.

Sinha, P., L. Brown, and S. Das. 1997. A re-examination of financial analysts' differential earnings forecast accuracy. *Contemporary Accounting Research* 14 (1): 1-42.

Soffer, L., S. Thiagarajan, and B. Walther. 2000. Earnings preannouncement strategies. *Review of Accounting Studies* 5 (1): 5-26.

Stickel, S. 1990. Predicting individual analyst earnings forecasts. *Journal of Accounting Research* 28 (2): 409-417.

Tevelow, R. 1971. How a security analyst uses the annual report. *Financial Executive* 39 (11): 18-21.

Thomas, J. 1993. Comments on 'Earnings forecasting research: Its implications for capital markets research' by L. Brown. *International Journal of Forecasting* 9: 325-330.

Womack, K. 1996. Do brokerage analysts' recommendations have investment value?" *The Journal of Finance* 51 (1): 137-167.

Zarnowitz, V. 1991. Has macro-forecasting failed? *The Cato Journal* 12 (1): 129-164.

Zmijewski, M. 1993. Comments on 'Earnings forecasting research: Its implications for capital markets research' by L. Brown." *International Journal of Forecasting* 9: 337-342.

**Figure 1a – Analyst Decision Process Schematic**

Panel A: Decision process schematic

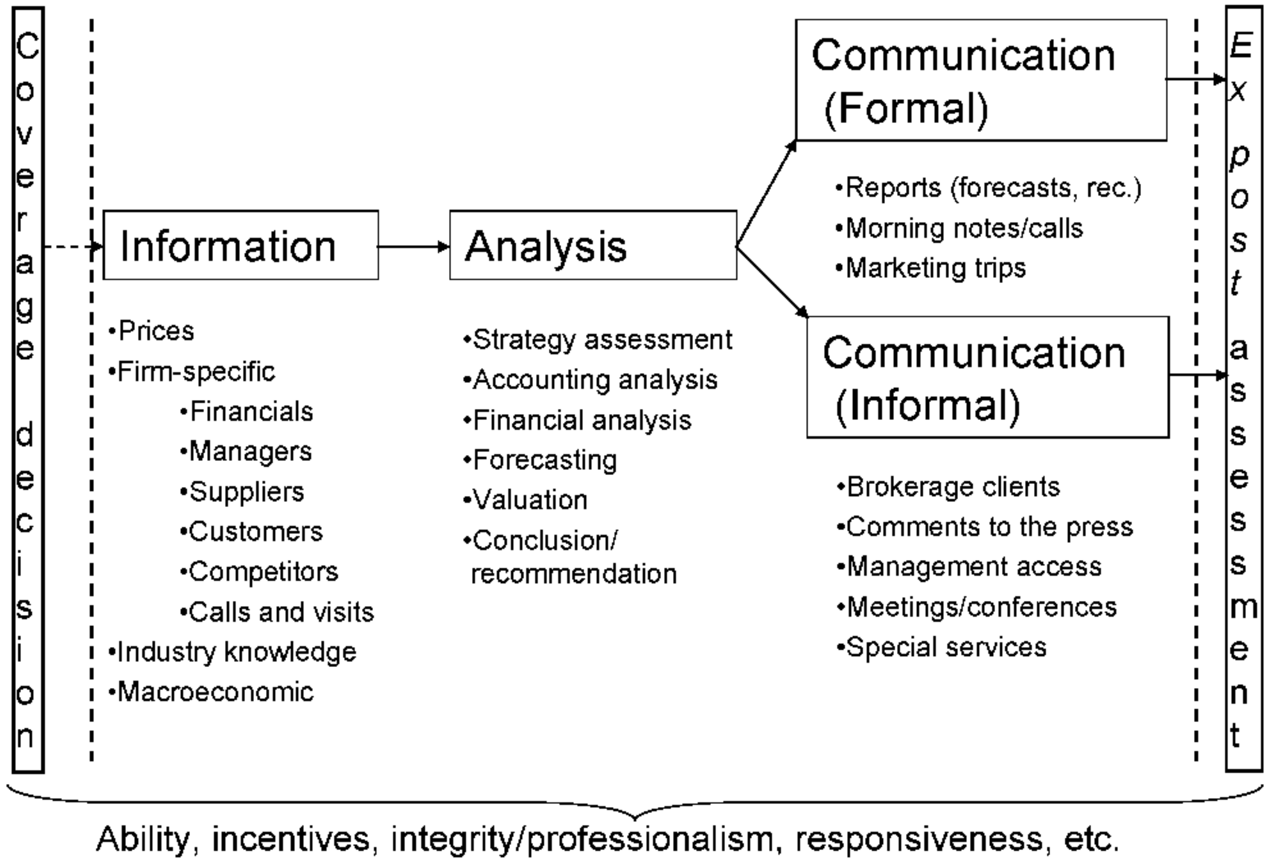

**Figure 1b – Analyst Decision Process Schematic (cont.)**

Panel A: Decision process schematic with most common research designs indicated

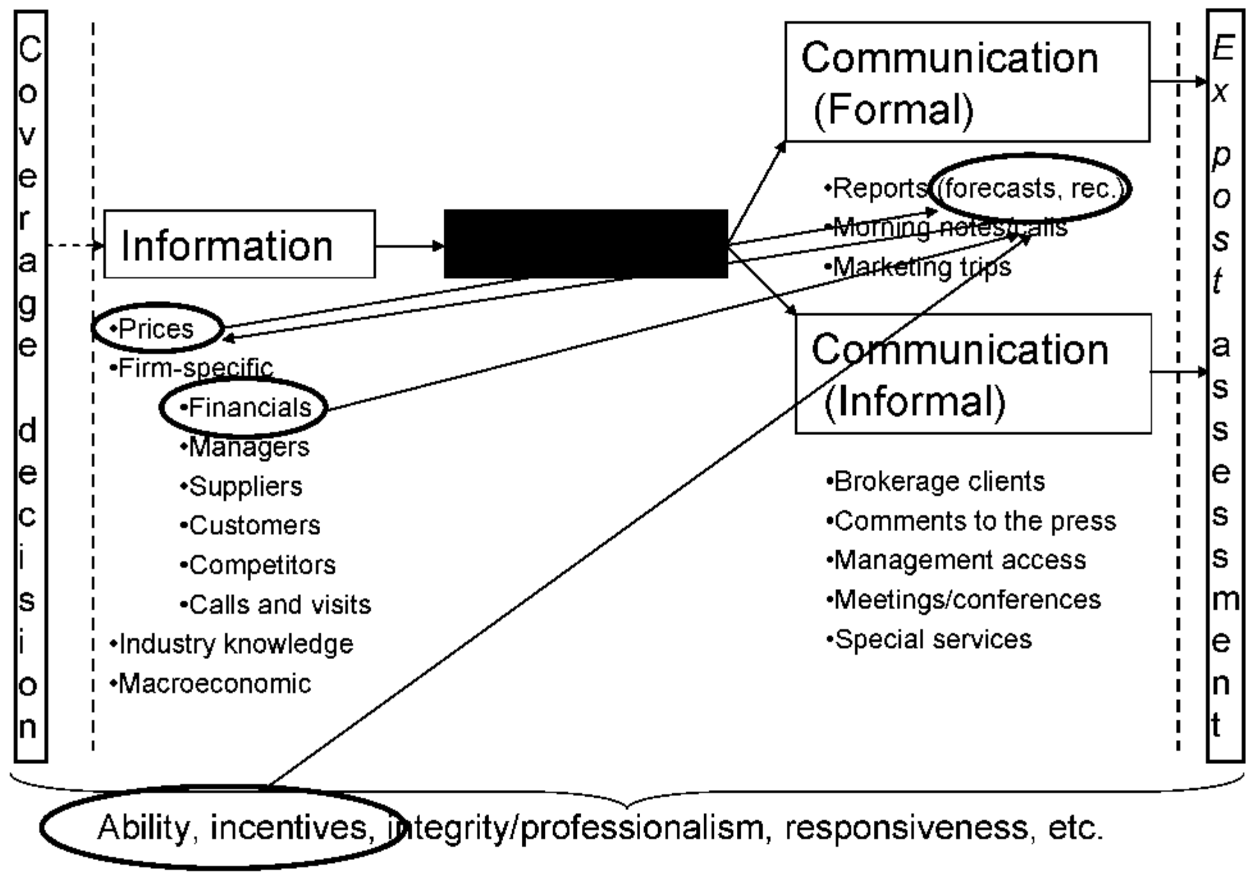

**Figure 2 – Timeline of Major Areas of Research 1968-2006**

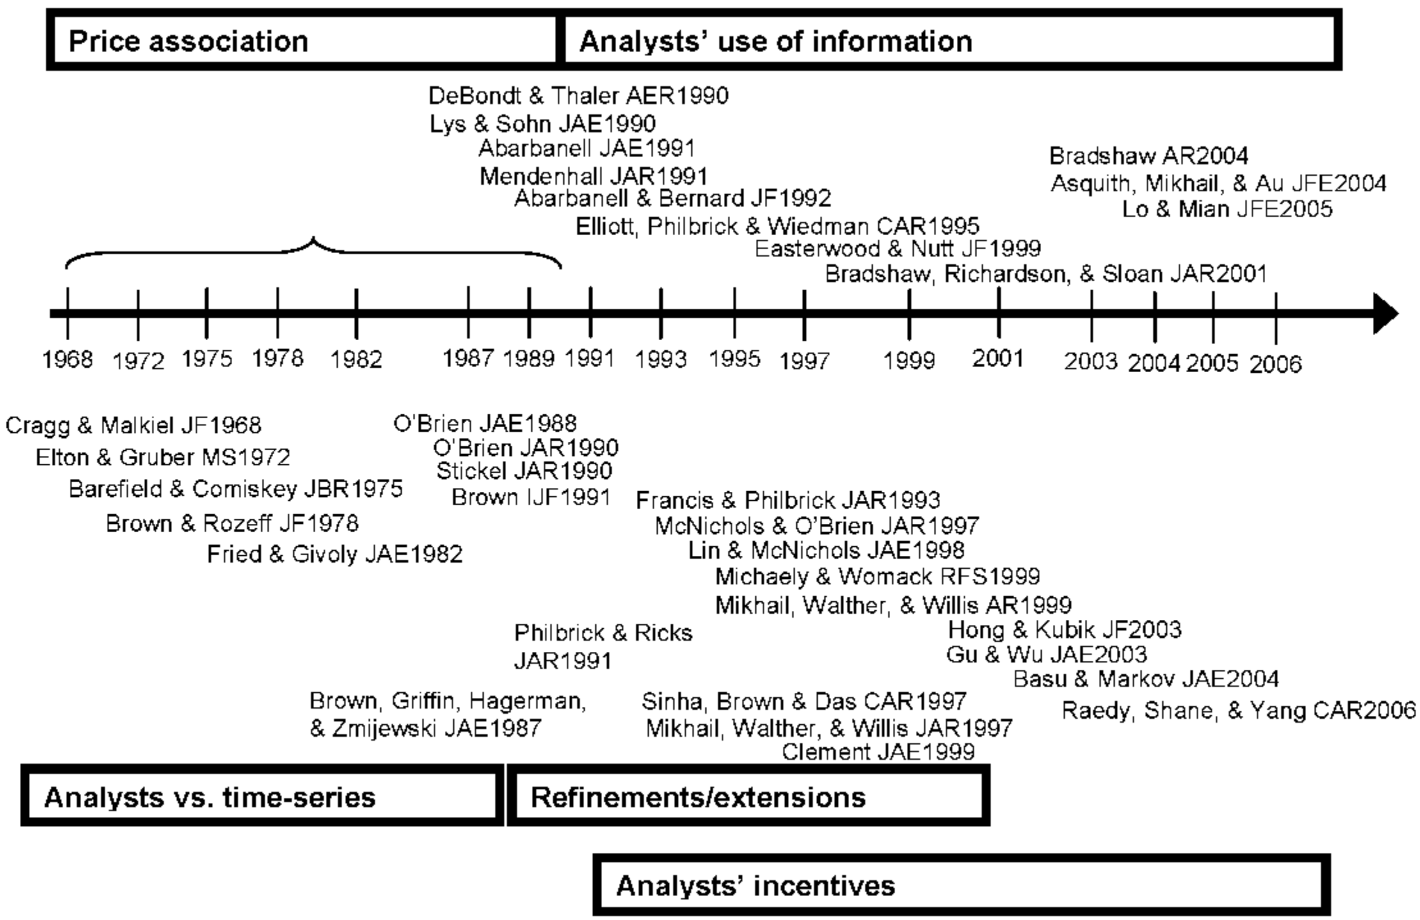

**Table 1 – Summary of Institutional Investor Ranking Surveys 1998-2005**

| | 1998 | 1999 | 2000 | 2001 | 2002 | 2003 | 2004 | 2005 |
|---|---|---|---|---|---|---|---|---|
| Industry knowledge | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 |
| Integrity/professionalism | | | | | 2 | 2 | 2 | 2 |
| Accessibility/responsiveness | | | | 2 | 3 | 3 | 3 | 3 |
| Management access | | | | 7 | 5 | 5 | 4 | 4 |
| Special services | 4 | 3 | 2 | 5 | 7 | 6 | 5 | 5 |
| Written reports | 3 | 2 | 4 | 6 | 8 | 7 | 7 | 6 |
| Timely calls and visits | | | | 4 | 4 | 4 | 6 | 7 |
| Communication skills | | | | | 10 | 9 | 8 | 8 |
| Financial models | | | 3 | 8 | 9 | 10 | 10 | 9 |
| Management of conflicts of interest | | | | 3 | 6 | 8 | 9 | 10 |
| Stock selection | 2 | 5 | 7 | 10 | 11 | 11 | 11 | 11 |
| Earnings estimates | 5 | 6 | 5 | 9 | 12 | 12 | 12 | 12 |
| Quality of sales force | 7 | 7 | 8 | 11 | 13 | 13 | 13 | |
| Market making | 8 | 8 | 9 | 12 | 14 | 14 | 14 | |
| Primary market services | | | 10 | | 15 | 15 | 15 | |
| Servicing | 6 | 4 | 6 | | | | | |

**Table 2 – Change in Ranked Characteristics, Institutional Investor Ranking Surveys 1998-2005**

| | Avg. rank change, 98-05 |
|---|---|
| (#2) Integrity/professionalism | 0.13 |
| (#3) Accessibility/responsiveness | 0.12 |
| Management access | 0.11 |
| Timely calls and visits | 0.07 |
| Communication skills | 0.06 |
| Financial models | 0.05 |
| Management of conflicts of interest | 0.04 |
| Special services | 0.01 |
| | |
| (#1) Industry knowledge | 0.00 |
| Primary market services | 0.00 |
| | |
| Market making | -0.02 |
| Written reports | -0.02 |
| Quality of sales force | -0.04* |
| Servicing | -0.05 |
| Earnings estimates | -0.06* |
| Stock selection | -0.10*** |

THE WALL STREET JOURNAL.

# Analysts: Still Coming Up Rosy --- Over-Optimism on Growth Rates Is Rampant, and the Estimates Help to Buoy Market's Valuation

*By Ken Brown*. **Wall Street Journal**. (Eastern edition).New York, N.Y.: Jan 27, 2003. pg. C.1



WALL STREET IS pretty downcast these days, what with a $1.5 billion settlement pending with regulators over stock-research conflicts, continuing layoffs at big securities firms and a stock market that is teetering yet again -- not to mention a cold snap that could freeze the thumbs of Blackberry users.

Yet stock analysts are unshaken in their optimistic, if delusional, belief that most of the companies they cover will have above-average, double-digit growth rates during the next several years. That is, of course, highly unlikely. Historically, corporate earnings have grown at about the same rate as the economy over time, and few expect the economy to grow at a double-digit rate any time soon.

But analysts refuse to bend to reality. Of the companies in the Standard & Poor's 500-stock index, analysts expect 345 of them to boost their earnings more than 10% a year during the next three to five years, and 123 companies to grow more than 15%, according to Multex, a stock-market-data firm.

"Hope springs eternal," says Mark Donovan, who manages Boston Partners Large Cap Value Fund. "You would have thought that, given what happened in the last three years, people would have given up the ghost. But in large measure they have not."

These overly optimistic growth estimates also show that, even with all the regulatory focus on too-bullish analysts allegedly influenced by their firms' investment-banking relationships, a lot of things haven't changed: Research remains rosy and many believe it always will.

In some ways, these high estimated growth rates underpin the market's current valuation, which remains pricey by historical standards. Investors expect to pay a higher price for stocks that are growing strongly. So if people realize these long-term growth-rate numbers are largely fictional, then a pillar of support for the market's valuation -- the S&P 500 currently trades at a price-to-earnings ratio of 18.5 based on 2002 earnings -- could go out of the stock market, sending prices lower.

The long-term growth figures come from the earnings estimates Wall Street analysts post for the companies they cover. Besides issuing buy and sell recommendations and predicting earnings during the next few quarters, analysts typically estimate how quickly the companies' earnings will grow during the next few years. Such long-term growth-rate numbers, which are imprecise by nature, give a hint of how analysts feel about companies' future prospects.

A long-term growth-rate number is often used by investors to determine whether a stock is cheap or expensive. Online auctioneer eBay Inc., for example, trades at a price-to-earnings ratio of 88 based on the past year's earnings. Some investors take solace in the fact that the company is expected to expand earnings 40% a year, but even with that growth, it would take until 2006 for the company's price-to-earnings ratio to fall to 22, assuming the stock price remained stalled at today's level.

These rosy figures come on top of three years of little or no growth for many companies. For example, Charles Schwab Corp. hasn't grown at all since 2000 as it has struggled with the stock-market collapse. But analysts, on average, still expect the company will expand its earnings 18% a year during the next several years. While that doesn't justify the company's price-to-earnings ratio of 33, it does give some hope to shareholders that the company one day indeed could resume its old growth rate.

Not surprisingly, the glow is rosiest in the technology sector. Of the 91 tech companies in the S&P 500, analysts expect 82 to grow faster than 10% a year, and 18 to grow better than 20% a year, meaning tech companies account for more than half of the index's 35 top growers.

To be sure, many of these companies could actually meet those growth expectations, if only because earnings have been in such a slump they are bound to rebound at some point. Analysts expect Schwab, for example, to earn 40 cents a share in 2003, up from the 29 cents it earned last year. If the analysts are right, that would be a healthy 38% jump in earnings.

But some also concede that their growth rates are optimistic. Guy Moszkowski, who covers Schwab for Salomon Smith Barney, and whose long-term growth estimate of 18% matches the consensus, concedes that this figure might be optimistic in the years after the expected short-term earnings pop. "If we can get enough of a recovery in the market that they can achieve that 40 cents in earnings, then they'll be on the way to establishing a kind of mid-teens growth track," he says. "But I think it's really hard to make the case they can do much better than that."

Mark Constant, who covers the company for Lehman Brothers and has a 15%-a-year growth estimate, also says the company probably won't reach his target. "I've always characterized it in print as an optimistic growth rate," he says.

If it were true that analysts were expecting a rebound following the current slump and ratcheting up their expectations accordingly, they might now be able to argue that they aren't being overly optimistic. The truth is, however, they have been growing increasingly pessimistic since the tech-stock bubble burst. Back in mid 2000, when earnings had been soaring for years, analysts were predicting that earnings for the S&P 500 would continue growing 15% a year, according to Morgan Stanley. Now, they are predicting 12% annual earnings growth for these same companies.

You can't blame analysts for everything,though. Companies themselves are guilty of being overly optimistic as well. "I think there's an immense amount of inertia in the system. That's the problem," says Steve Galbraith, Morgan Stanley's chief investment strategist. "One of the things people are struggling with are creative ways of reducing your guidance without reducing your guidance."

The problem, he adds, is that many companies set their growth expectations a decade ago, when interest rates and inflation were higher than today. Growth rates are measured in nominal terms, meaning inflation gives them a boost. With virtually no inflation and interest rates near zero, it is harder for companies to post double-digit growth. "I do think this is something that corporate America broadly is wrestling with: How do we ratchet down expectations that we set 10 years ago when things were different?" he says.

The danger comes from companies that can't face the reality that their growth has slowed. "Where I think clients should get concerned is where a company is claiming they're a 15% grower and they're setting their capital expenditures accordingly," Mr. Galbraith says. If the market is pricing in that level of growth, then the company will likely keep investing in itself in an attempt to keep returns high. The danger of that: Companies could be throwing away capital that could be given back to investors in the form of dividends or share buybacks.

Every chief financial officer who took Corporate Finance 101 knows that the bigger the portion of earnings a company reinvests in its business, the faster it conceivably can grow. Sending cash out to investors reduces the amount the company can invest in itself, ultimately lowering its potential growth rate.

But there are signs -- including Microsoft Corp.'s plan to pay a dividend -- that executives are starting to realize that reinvesting all their excess cash in their own business might not produce the highest returns. "It hasn't gotten quite that far, but I think it's going to get there," says Jeff van Harte, who manages Transamerica Premier Equity fund. "It just takes a long time to change attitudes. Some companies are forever lost."

---

BUSINESS INSIDER

# Interest rate forecasters are shockingly wrong almost all of the time

AKIN OYEDELE
JUL. 8, 2015, 8:25 AM


Wikimedia Commons

Most interest rate forecasters are wrong most of the time.

Very wrong.

The chart below is from Jeff Gundlach's presentation on Tuesday, comparing the US 10-year yield to median economist forecasts over the past five years.

The black line is the 10-year yield, and the colored lines are the paths that economists thought rates would take.

Clearly, these forecasters were wrong most of the time, as there were only a few instances of convergence between both lines.

In 2012, forecasters were hugely bleak about the economy, and thought that interest rates would collapse the whole year. Rates ended the year higher than where they started.

Last year was particularly bad, when strategists became too optimistic that the Federal Reserve would hike rates.

This year, forecasters again thought rates would rise and as rates fell, so did those forecasts, which have now converged with interest rates.

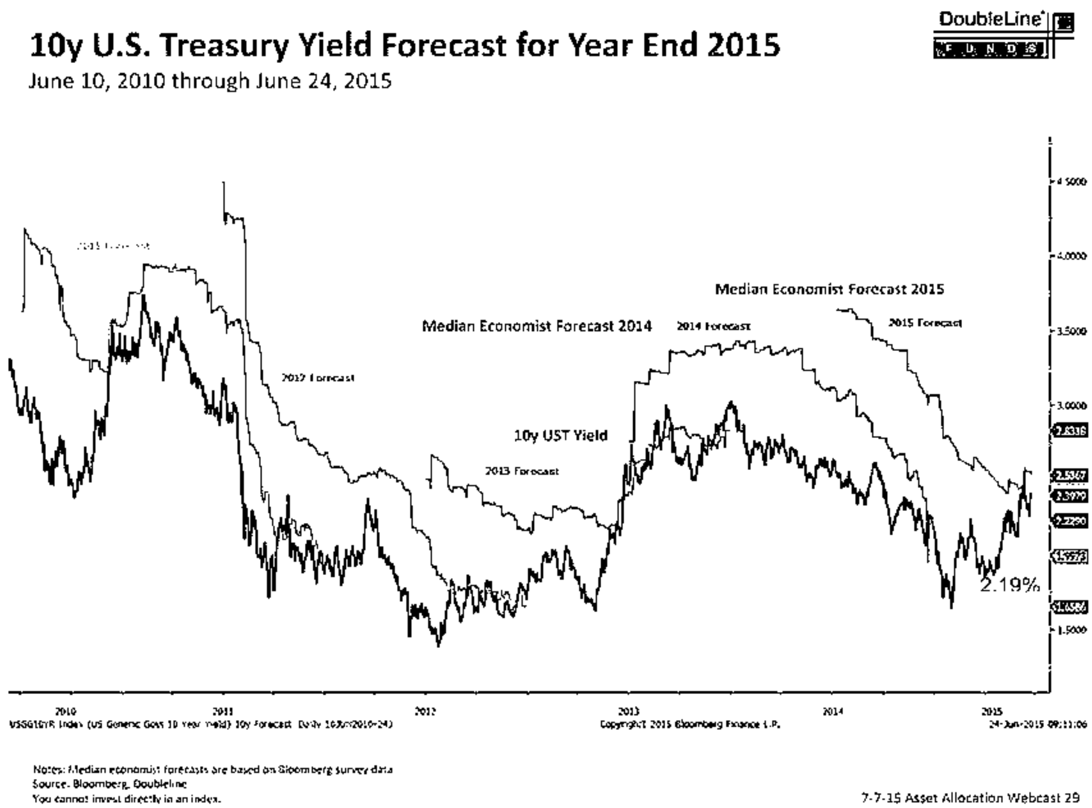

*Doubleline Funds*

## NOW WATCH: Someone figured out the purpose of the extra shoelace hole on your running shoes — and it will blow your mind

More: Interest Rates Jeff Gundlach Forecasting

- 
- 
- 

## Callan's 10-year assumptions by asset class

Explore our long-term outlook for return and risk for all of the major asset classes. In this chart, you can view those assumptions by asset class type, such as equities or fixed income. Simply hover over the various bars in the chart to see the relevant index.

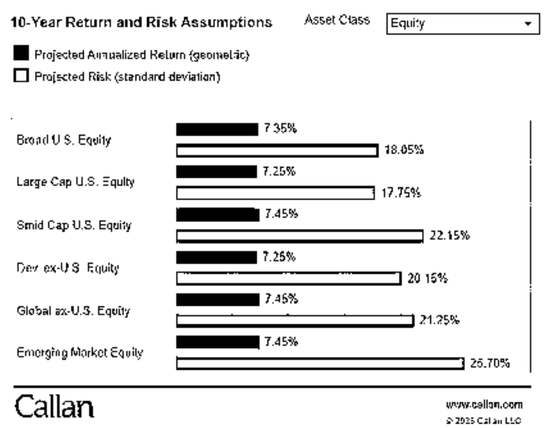


https://www.callan.com/capital-markets-assumptions/

7.25% - 3.5% = 3.75%

**How Callan's 10-Year Return Assumptions Have Changed Since 2003**

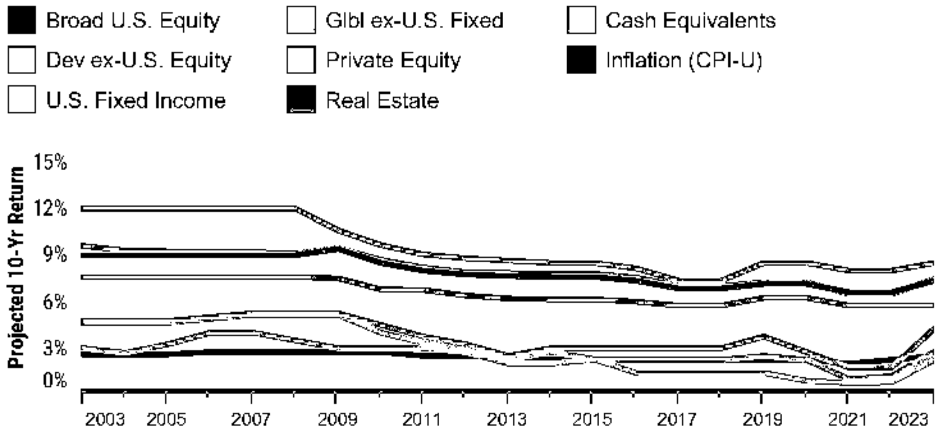


Callan



https://www.callan.com/capital-markets-assumptions/

**Forecasting US Equity Returns in the 21st Century**

John Y. Campbell, Harvard University
July 2001

What returns should investors expect the US stock market to deliver on average during the next century? Does the experience of the last century provide a reliable guide to the future? In this short note I first discuss alternative methodologies for forecasting average future equity returns, then discuss current market conditions, and finally draw conclusions for long-term return forecasts. Throughout I work in real, that is inflation-adjusted, terms.

## I. Methods for forecasting returns

### 1. Average past returns

Perhaps the simplest way to forecast future returns is to use some average of past returns. Very naturally, this method has been favored by many investors and analysts. However there are several difficulties with it.

a) *Geometric average or arithmetic average?* The geometric average return is the cumulative past return on US equities, annualized. Siegel (1998) studies long-term historical data on value-weighted US share indexes. He reports a geometric average of 7.0% over two different sample periods, 1802 1997 and 1871 1997. The arithmetic average return is the average of one-year past returns on US equities. It is considerably higher than the geometric average return, 8.5% over 1802 1997 and 8.7% over 1871 1997.[1]

When returns are serially uncorrelated, the arithmetic average represents the best forecast of future return in any randomly selected future year. For long holding periods, the best forecast is the arithmetic average compounded up appropriately. If one is making a 75-year forecast, for example, one should forecast a cumulative return of $1.085^{75}$ based on 1802 1997 data.

When returns are negatively serially correlated, however, the arithmetic average is not necessarily superior as a forecast of long-term future returns. To understand this, consider an extreme example in which prices alternate deterministically between 100 and 150. The return is 50% when prices rise, and -33% when prices fall. Over any even number of periods, the geometric average return is zero, but the arithmetic average return is 8.5%. In this case the arithmetic average return is misleading because it fails to take account of the fact that high returns always multiply a low initial price of 100, while low returns always multiply a high initial price of 150. The geometric average is a better indication of long-term future

[1] When returns are lognormally distributed, the difference between the two averages is approximately one-half the variance of returns. Since stock returns have an annual standard deviation of about 18% over these long periods, the predicted difference is $0.18^2/2$ $0.016$ or 1.6%. This closely matches the difference in the data.

prospects in this example.[2]

This point is not just a theoretical curiosity, because in the historical data summarized by Siegel, there is strong evidence that the stock market is mean-reverting. That is, periods of high returns tend to be followed by periods of lower returns. This suggests that the arithmetic average return probably overstates expected future returns over long periods.

b) *Returns are very noisy.* The randomness in stock returns is extreme. With an annual standard deviation of real return of 18%, and 100 years of past data, a single year's stock return that is only one standard deviation above average increases the average return by 18 basis points. A lucky year that is two standard deviations above average increases the average return by 36 basis points. Even when a century or more of past data is used, forecasts based on historical average returns are likely to change substantially from one year to the next.

c) *Realized returns rise when expected returns fall.* To the extent that expected future equity returns are not constant, but change over time, they can have perverse effects on realized returns. Suppose for example that investors become more risk-tolerant and reduce the future return that they demand from equities. If expected future cash flows are unchanged, this drives up prices and realized returns. Thus an estimate of future returns based on average past realized returns will tend to increase just as expected future returns are declining.

Something like this probably occurred in the late 1990's. A single good year can have a major effect on historical average returns, and several successive good years have an even larger effect. But it would be a mistake to react to the spectacular returns of 1995–99 by increasing estimates of 21st Century returns.

d) *Unpalatable implications.* Fama and French (2000) point out that average past US stock returns are so high that they exceed estimates of the return to equity (ROE) calculated for US corporations from accounting data. Thus if one uses average past stock returns to estimate the cost of capital, the implication is that US corporate investments have destroyed value; corporations should instead have been paying all their earnings out to stockholders. This conclusion is so hard to believe that it further undermines confidence in the average-return methodology.

One variation of the average-past-returns approach is worth discussing. One might take the view that average past equity returns in other countries provide relevant evidence about US equity returns. Standard international data from Morgan Stanley Capital International, available since the early 1970's, show that equity returns in most other industrialized countries have been about as high as those in the US. The exceptions are the heavily commodity-dependent markets of Australia and Canada, and the very small Italian market (Campbell 1999). Jorion and Goetzmann (1999) argue that other countries' returns were

[2]One crude way to handle this problem is to measure the annualized variance of returns over a period such as 20 years that is long enough for returns to be approximately serially uncorrelated, and then to adjust the geometric average up by one-half the annualized 20-year variance as would be appropriate if returns are lognormally distributed. Campbell and Viceira (2001, Figure 4.2) report an annualized 20-year standard deviation of about 14% in long-term annual US data, which would imply an adjustment of $0.14^2/2 = 0.010$ or 1.0%.

lower than US returns in the early 20th Century, but this conclusion appears to be sensitive to their omission of the dividend component of return (Dimson, Marsh, and Staunton 2000). Thus the use of international data does not change the basic message that the equity market has delivered high average returns in the past.

### 2. Valuation ratios

An alternative approach is to use valuation ratios ratios of stock prices to accounting measures of value such as dividends or earnings to forecast future returns. In a model with constant valuation ratios and growth rates, the famous Gordon growth model says that the dividend-price ratio

$$\frac{D}{P} = R - G, \tag{1}$$

where $R$ is the discount rate or expected equity return, and $G$ is the growth rate of dividends (equal to the growth rate of prices when the valuation ratio is constant). This formula can be applied either to price per share and conventional dividends per share, or to the total value of the firm and total cash paid out by the firm (including share repurchases). A less well-known but just as useful formula says that in steady state, where earnings growth comes from reinvestment of retained earnings which earn an accounting ROE equal to the discount rate $R$,

$$\frac{E}{P} = R. \tag{2}$$

Over long periods of time summarized by Siegel (1998), these formulas give results consistent with average realized returns. Over the period 1802 1997, for example, the average dividend-price ratio was 5.4% while the geometric average growth rate of prices was 1.6%. These numbers add to the geometric average return of 7.0%. Over the period 1871 1997 the average dividend-price ratio was 4.9% while the geometric average growth rate of prices was 2.1%, again adding to 7.0%. Similarly, Campbell and Shiller (2001) report that the average P/E ratio for S&P500 shares over the period 1872-2000 was 14.5. The reciprocal of this is 6.9%, consistent with average realized returns.

When valuation ratios and growth rates change over time, these formulas are no longer exactly correct. Campbell and Shiller (1988) and Vuolteenaho (2000) derive dynamic versions of the formulas that can be used in this context. Campbell and Shiller show, for example, that the log dividend-price ratio is a discounted sum of expected future discount rates, less a discounted sum of expected future dividend growth rates. In this note I will work with the simpler deterministic formulas.

## II. Current market conditions

Current valuation ratios are wildly different from historical averages, reflecting the unprecedented bull market of the last 20 years, and particularly the late 1990's. The attached figure, taken from Campbell and Shiller (2001), illustrates this point. The bottom left panel shows the dividend-price ratio $D/P$ in January of each year from 1872–2000. The long-term historical average is 4.7%, but $D/P$ has fallen dramatically since 1982 to about 1.2% in January 2000 (and 1.4% today).

The dividend-price ratio may have fallen in part because of shifts in corporate financial policy. An increased tendency for firms to repurchase shares rather than pay dividends increases the growth rate of dividends per share, by shrinking the number of shares. Thus it increases $G$ in the Gordon growth formula and reduces conventionally measured $D/P$. One way to correct for this is to add repurchases to conventional dividends. Recent estimates of this effect by Liang and Sharpe (1999) suggest that it may be an upward adjustment of 75 to 100 basis points, and more in some years. Of course, this is not nearly sufficient to explain the recent decline in $D/P$.

Alternatively, one can look at the price-earnings ratio. The top left panel of the figure shows $P/E$ over the same period. This has been high in recent years, but there are a number of earlier peaks that are comparable. Close inspection of these peaks shows that they often occur in years such as 1992, 1934, and 1922 when recessions caused temporary drops in (previous-year) earnings. To smooth out this effect, Campbell and Shiller (2001), following Graham and Dodd (1934), advocate averaging earnings over 10 years. The price-averaged earnings ratio is illustrated in the top right panel of the figure. This peaked at 45 in January 2000; the previous peak was 28 in 1929. The decline in the S&P500 since January 2000 has only brought the ratio down to the mid-30's, still higher than any level seen before the late 1990's.

The final panel in the figure, on the bottom right, shows the ratio of current to 10-year average earnings. This ratio has been high in recent years, reflecting robust earnings growth during the 1990's, but it is not unprecedentedly high. The really unusual feature of the recent stock market is the level of prices, not the growth of earnings.

### III. Implications for future returns

The implications of current valuations for future returns depend on whether the market has reached a new steady state, in which current valuations will persist, or whether these valuations are the result of some transitory phenomenon.

If current valuations represent a new steady state, then they imply a substantial decline in the equity returns that can be expected in the future. Using Campbell and Shiller's (2001) data, the unadjusted dividend-price ratio has declined by 3.3 percentage points from the historical average. Even adjusting for share repurchases, the decline is at least 2.3 percentage points. Assuming constant long-term growth of the economy, this would imply that the geometric average return on equity is no longer 7%, but 3.7% or at most 4.7%. Looking at the price-averaged earnings ratio, adjusting for the typical ratio of current to averaged earnings, gives an even lower estimate. Current earnings are normally 1.12 times averaged earnings; $1.12/35 = 0.032$, implying a 3.2% return forecast. These forecasts allow for only a very modest equity premium relative to the yield on long-term inflation-indexed bonds, currently about 3.5%, or the 3% safe real return assumed recently by the Trustees.

If current valuations are transitory, then it matters critically what happens to restore traditional valuation ratios. One possibility is that earnings and dividends are below their long-run trend levels; rapid earnings and dividend growth will restore traditional valuations without any declines in equity returns below historical levels. While this is always a possibility, Campbell and Shiller (2001) show that it would be historically unprecedented. The US stock market has an extremely poor record of predicting future earnings and dividend growth. Historically stock prices have increased relative to earnings during decades of rapid earnings growth, such as the 1920's, 1960's, or 1990's, as if the stock market anticipates that rapid earnings growth will continue in the next decade. However there is no systematic tendency for a profitable decade to be followed by a second profitable decade; the 1920's, for example, were followed by the 1930's and the 1960's by the 1970's. Thus stock market optimism often fails to be justified by subsequent earnings growth.[3]

A second possibility is that stock prices will decline or stagnate until traditional valuations are restored. This has occurred at various times in the past after periods of unusually high stock prices, notably the 1900's and 1910's, the 1930's, and the 1970's. This would imply extremely low and perhaps even negative returns during the adjustment period, and then higher returns afterwards.

The unprecedented nature of recent stock market behavior makes it impossible to base forecasts on historical patterns alone. One must also form a view about what happened to drive stock prices up during the 1980's and particularly the 1990's. One view is that there has been a structural decline in the equity premium, driven either by the correction of mistaken perceptions of risk (aided perhaps by the work of economists on the equity premium puzzle), or by the reduction of barriers to participation and diversification by small investors.[4] Economists such as McGrattan and Prescott (2001) and Jagannathan, McGrattan, and Scherbina (2001) argue that the structural equity premium is now close to zero, consistent with theoretical models in which investors effectively share risks and have modest risk aversion, and consistent with the view that the US market has reached a new steady state.

An alternative view is that the equity premium has declined only temporarily, either because investors irrationally overreacted to positive fundamental news in the 1990's (Shiller 2000), or because the strong economy made investors more tolerant of risk.[5] On this view the equity premium will return to historical levels, implying extremely poor near-term returns and higher returns in the more distant future after traditional valuations have been restored.

It is too soon to tell which of these views is correct, and I believe it is sensible to put some weight on each of them. That is, I expect valuation ratios to return part way but not

[3]Vuolteenaho (2000) notes, however, that US corporations were unusually profitable in the late 1990's and that profitability has some predictive power for future earnings growth.

[4]Heaton and Lucas (1999) model barriers of this sort. It is hard to get large effects of increased participation on stock prices unless initial participation levels are extremely low. Furthermore, one must keep in mind that what matters for pricing is the wealth-weighted participation rate, that is, the probability that a randomly selected dollar of wealth is held by an individual who can participate in the market. This is higher than the equal-weighted participation rate, the probability that a randomly selected individual can participate.

[5]Campbell and Cochrane (1999) present a model in which investors judge their well-being by their consumption relative to a recent average of past aggregate consumption. In this model investors are more risk-tolerant when consumption grows rapidly and they have a "cushion of comfort" relative to their minimum expectations. The Campbell-Cochrane model fits past cyclical variations in the stock market, which will likely continue in the future, but it is hard to explain the extreme recent movements using this model.

fully to traditional levels.[6] A rough guess for the long term, after the adjustment process is complete, might be a geometric average equity return of 5% to 5.5% or an arithmetic average return of 6.5% to 7%.

If equity returns are indeed lower on average in the future, it is likely that short-term and long-term real interest rates will be somewhat higher. That is, the total return to the corporate capital stock is determined primarily by the production side of the economy and by national saving and international capital flows; the division of total return between riskier and safer assets is determined primarily by investor attitudes towards risk. Reduced risk aversion then reduces the equity premium both by driving down the equity return and by driving up the riskless interest rate. The yield on long-term inflation-indexed Treasury securities (TIPS) is about 3.5%, while short-term real interest rates have recently averaged about 3%. Thus 3% to 3.5% would be a reasonable guess for safe real interest rates in the future, implying a long-run average equity premium of 1.5% to 2.5% in geometric terms or about 3% to 4% in arithmetic terms.

Finally, I note that it is tricky to use these numbers appropriately in policy evaluation. Average equity returns should never be used in base-case calculations without showing alternative calculations to reflect the possibilities that realized returns will be higher or lower than average. These calculations should include an alternative in which equities underperform Treasury bills. Even if the probability of underperformance is small over a long holding period, it cannot be zero or the stock market would be offering an arbitrage opportunity or "free lunch". Equally important, the bad states of the world in which underperformance occurs are heavily weighted by risk-averse investors. Thus policy evaluation should use a broad range of returns to reflect the uncertainty about long-run stock market performance.

[6] This compromise view also implies that negative serial correlation, or mean-reversion, is likely to remain a characteristic of stock returns in the 21st Century.

## Bibliography


Campbell, John Y. and John H. Cochrane, 1999, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior", *Journal of Political Economy* 107, 205–251.

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

Campbell, John Y. and Robert J. Shiller, 2001, "Valuation Ratios and the Long-Run Stock Market Outlook: An Update", NBER Working Paper No. 8221.

Campbell, John Y. and Luis M. Viceira, 2001, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, forthcoming Oxford University Press, New York, NY.

Dimson, Elroy, Paul Marsh, and Mike Staunton, 2000, *The Millennium Book: A Century of Investment Returns*, ABN Amro and London Business School, London.

Fama, Eugene and Kenneth R. French, 2000, "The Equity Premium", unpublished paper, University of Chicago and MIT.

Graham, Benjamin and David Dodd, 1934, *Security Analysis*, McGraw-Hill, New York, NY.

Heaton, John and Deborah Lucas, 1999, "Stock Prices and Fundamentals", *NBER Macroeconomics Annual* 213–242, MIT Press, Cambridge, MA.

Jagannathan, Ravi, Ellen R. McGrattan, and Anna Scherbina, 2001, "The Declining US Equity Premium", NBER Working Paper No. 8172.

Jorion, Philippe and William N. Goetzmann, 1999, "Global Stock Markets in the Twentieth Century", *Journal of Finance* 54, 953–980.

Liang, Nellie and Steven A. Sharpe, 1999, "Share Repurchases and Employee Stock Options and their Implications for S&P500 Share Retirements and Expected Returns", unpublished paper, Board of Governors of the Federal Reserve System.

McGrattan, Ellen R., and Edward C. Prescott, 2001, "Is the Stock Market Overvalued?", NBER Working Paper No. 8077.

Shiller, Robert J., 2000, *Irrational Exuberance*, Princeton University Press, Princeton, NJ.

Siegel, Jeremy, 1998, *Stocks for the Long Run*, 2nd ed., McGraw-Hill, New York, NY.

Vuolteenaho, Tuomo, 2000, "Understanding the Aggregate Book-to-Market Ratio and its Implications to Current Equity-Premium Expectations", unpublished paper, Harvard University.

Figure 4. **S&P Composite Stock Data, January Values 1872-1997**

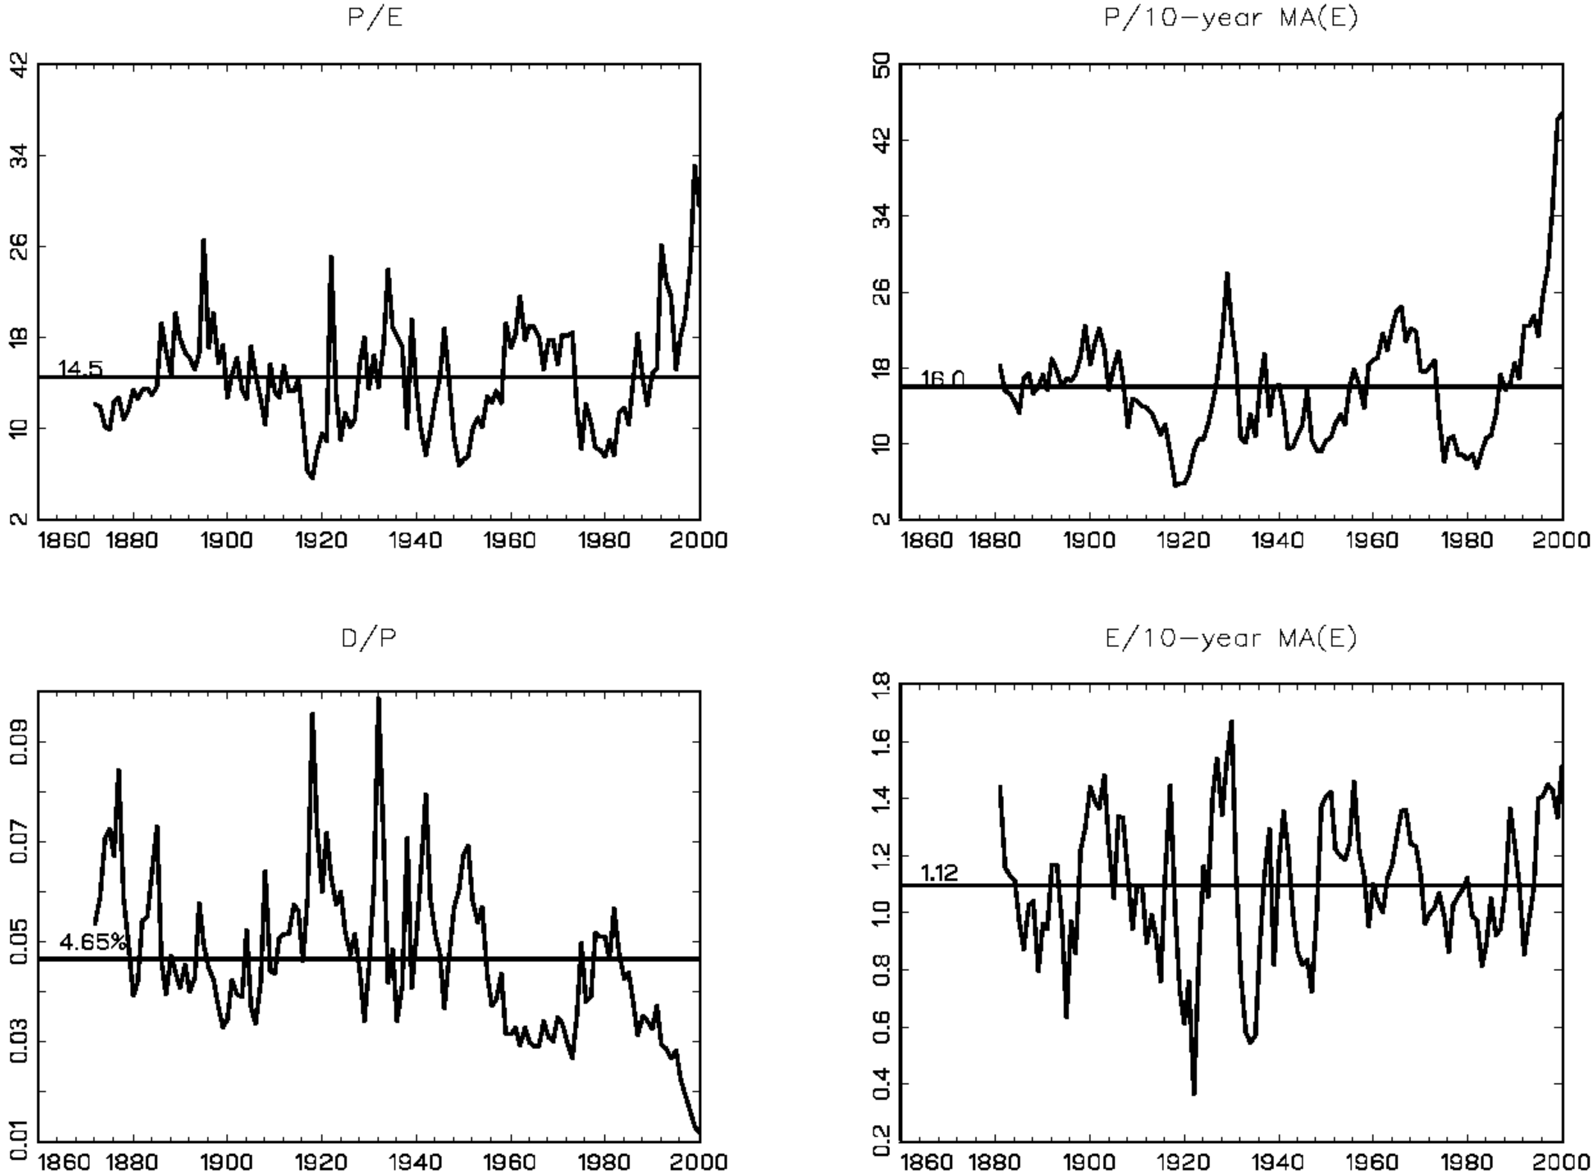

# Viewpoint: Estimating the equity premium

John Y. Campbell *Department of Economics, Harvard University*

*Abstract.* Finance theory restricts the time-series behaviour of valuation ratios and links the cross-section of stock prices to the level of the equity premium. This can be used to strengthen the evidence for predictability in stock returns. Steady-state valuation models are useful predictors of stock returns, given the persistence in valuation ratios. A steady-state approach suggests that the world geometric average equity premium fell considerably in the late twentieth century, rose modestly in the early years of the twenty-first century, and was almost 4% at the end of March 2007. JEL classification: G12

*Evaluer la prime des actions par rapport aux obligations.* La théorie financière contraint le comportement diachronique des ratios de valorisation et relie transversalement les prix des actions au niveau de prime des actions sur les obligations. Voilà qui peut être utilisé pour renforcer la prédictibilité des rendements sur les actions. Les modèles de valorisation en régime permanent sont des prédicteurs utiles des rendements sur les actions, compte tenu du caractère stable des ratios de valorisation. Une approche en termes de régime permanent suggère que la moyenne géométrique mondiale de la prime des actions sur les obligations a chuté considérablement à la fin du 20e siècle, qu'elle a été modestement en hausse dans les premières années du 21e siècle, et qu'elle était à presque 4% à la fin de mars 2007.

The author is also affiliated with Arrowstreet Capital, LP. and NBER. This paper was presented in June 2007 as a State of the Art lecture at the Canadian Economics Association annual meeting at Dalhousie University in Halifax, Nova Scotia. A precursor was presented in January 2007 to the D-CAF Conference on Return Predictability at Copenhagen Business School. I am grateful to participants at both conferences, to John Cochrane, Jon Lewellen, Lubos Pastor, Ivo Welch, and Jeff Wurgler, and particularly to Angelo Melino for their thoughtful comments; to Bob Shiller, Moto Yogo, and my colleagues at Arrowstreet Capital. Sam Thompson and Tuomo Vuolteenaho, for joint research and many conversations on this subject; and to Alex Ogan, also of Arrowstreet Capital, for his able assistance with the data illustrated in figures 1 through 5.
Email: john_campbell@harvard.edu.

Canadian Journal of Economics / Revue canadienne d'Economique, Vol. 41, No. 1
February / février 2008. Printed in Canada / Imprimé au Canada

0008-4085 / 08 / 1–21 / 

## 1. Introduction

What return should investors expect the stock market to deliver, above the interest rate on a safe short-term investment? In other words, what is a reasonable estimate of the equity premium?

This question is a basic one for investors who must decide how to allocate their portfolios to safe and risky assets. In the academic world, it has for over three decades played a central role in the development of asset pricing theory and financial econometrics. In the 1960s and 1970s, the efficient market hypothesis was interpreted to mean that the true equity premium was a constant. Investors might update their estimates of the equity premium as more data became available, but eventually these estimates should converge to the truth. This viewpoint was associated with the use of historical average excess stock returns to forecast future returns.

In the early 1980s, a number of researchers reported evidence that excess stock returns could be predicted by regressing them on lagged financial variables. In particular, valuation ratios that divide accounting measures of cash flow by market valuations, such as the dividend-price ratio, earnings-price ratio, or smoothed earnings-price ratio, appeared to predict returns. Value-oriented investors in the tradition of Graham and Dodd (1934) had always asserted that high valuation ratios are an indication of an undervalued stock market and should predict high subsequent returns, but these ideas did not carry much weight in the academic literature until authors such as Rozeff (1984), Fama and French (1988), and Campbell and Shiller (1988a,b) found that valuation ratios are positively correlated with subsequent returns and that the implied predictability of returns is substantial at longer horizons. Around the same time, several papers pointed out that yields on short- and long-term Treasury and corporate bonds are correlated with subsequent stock returns (Fama and Schwert 1977; Keim and Stambaugh 1986; Campbell 1987; Fama and French 1989).

These results suggested that the equity premium is not a constant number that can be estimated ever more precisely, but an unknown state variable whose value must be inferred at each point in time on the basis of observable data. Meanwhile, research in asset pricing theory made financial economists more comfortable with the idea that the equity premium can change over time even in an efficient market with rational investors, so that a time-varying equity premium does not necessarily require abandonment of the traditional paradigm of financial economics for a behavioural or inefficient-markets alternative. Campbell and Cochrane (1999), for example, showed that rational investors with habit formation preferences might become more averse to volatility in consumption and wealth, driving up the equilibrium equity premium, when the economy is weak.

During the 1990s, research continued on regressions predicting stock returns from valuation ratios (Kothari and Shanken 1997; Lamont 1998; Pontiff and Schall 1998) and interest rates (Hodrick 1992). However the 1990s also saw challenges to the new view that valuation ratios predict stock returns.

A first challenge came from financial econometricians, who began to express concern that the apparent predictability of stock returns might be spurious. Many of the predictor variables in the literature are highly persistent: Nelson and Kim (1993) and Stambaugh (1999) pointed out that persistence leads to biased coefficients in predictive regressions if innovations in the predictor variable are correlated with returns (as is strongly the case for valuation ratios, although not for interest rates). Under the same conditions the standard $t$-test for predictability has incorrect size (Cavanagh, Elliott, and Stock 1995). These problems are exacerbated if researchers are data mining, considering large numbers of variables and reporting only those results that are apparently statistically significant (Foster, Smith, and Whaley 1997; Ferson, Sarkissian, and Simin 2003). An active recent literature discusses alternative econometric methods for correcting the bias and conducting valid inference (Cavanagh, Elliott, and Stock 1995; Lewellen 2004; Torous, Valkanov, and Yan 2004; Campbell and Yogo 2006; Jansson and Moreira 2006; Polk, Thompson, and Vuolteenaho 2006; Ang and Bekaert 2007; Cochrane 2007).

A second challenge was posed by financial history. In the late 1990s valuation ratios were extraordinarily low, so regression forecasts of the equity premium became negative (Campbell and Shiller 1998). Yet stock returns continued to be high until after the turn of the millennium. Data from these years were sufficiently informative to weaken the statistical evidence for stock return predictability. Although low returns in the early 2000s have partially restored this evidence, Goyal and Welch (2003, 2007) and Butler, Grullon, and Weston (2005) have argued that overall, the out-of-sample forecasting power of valuation ratios is often worse than that of a traditional model predicting the equity premium using only the historical average of past stock returns.

The ultimate test of any predictive model is its out-of-sample performance. My personal experience using regression models to forecast stock returns in the late 1990s was humbling, although these models were partially vindicated by the stock market decline of the early 2000s. The lesson I draw from this experience is that one is more likely to predict stock returns successfully if one uses finance theory to reduce the number of parameters that must be freely estimated from the data and to restrict estimates of the equity premium to a reasonable range.

In the next section of this paper I show how finance theory can be used if one believes that valuation ratios, in particular the dividend-price ratio, are stationary around a constant mean. Even under stationarity, the persistence of valuation ratios has led researchers to concentrate on situations where valuation ratios have a root that is close to unity. In section 3 I discuss the limiting case where one believes that the dividend-price ratio follows a geometric random walk. I show that this case allows an even larger role for theory: it implies that one should forecast returns by adding a growth estimate to the dividend-price ratio, in the manner of the classic Gordon growth model. I argue that this approach has historically generated successful out-of-sample forecasts and is likely to do so in the future as well. In section 4 I apply this methodology to estimate the current equity premium for Canada, for the U.S., and for the world stock market as a whole. In section 5 I briefly discuss how finance theory can be used to predict the equity premium from the cross-section of stock prices. Section 6 concludes.

## 2. Regression-based return prediction with a stationary dividend-price ratio

When the dividend-price ratio is stationary, a basic tool for analysing stock returns is the loglinear approximate relation derived by Campbell and Shiller (1988a). This relation says that the log stock return $r_{t+1}$, the log stock price $p_t$, and the log dividend $d_t$ approximately satisfy

$$\begin{aligned} r_{t+1} &= k + \rho p_{t+1} + (1-\rho)d_{t+1} - p_t \\ &= k + (d_t - p_t) + \Delta d_{t+1} - \rho(d_{t+1} - p_{t+1}), \end{aligned} \tag{1}$$

where $\rho$ is a coefficient of loglinearization equal to the reciprocal of one plus the steady-state level of the dividend-price ratio. Thus $\rho$ is slightly smaller than one; for annual U.S. data, $\rho = 0.96$ is a reasonable value, given an average dividend-price ratio in the late twentieth century of about 4% or 0.04 in levels. This equation says that proportional changes in stock prices have a larger effect on returns than equal proportional changes in dividends, because the level of dividends is small relative to the level of prices.

Equation (1) is a difference equation for the log dividend-price ratio. Solving it forward, imposing a condition that there are no explosive bubbles in stock prices, and taking expectations at time $t$ allows us to interpret the dividend-price ratio as

$$d_t - p_t = \frac{k}{1-\rho} + E_t \sum_{j=0}^{\infty} \rho^j [r_{t+1+j} - \Delta d_{t+1+j}]. \tag{2}$$

This formula delivers a number of insights. First, it helps to motivate regressions of stock returns on the log dividend-price ratio. The ratio is a linear combination of discounted expectations of future stock returns and dividend growth. If dividend growth is not too predictable (and there is little direct evidence for long-term dividend predictability in U.S. data), and if the dynamics of discount rates are such that short- and long-term expected stock returns are highly correlated, then the log dividend-price ratio should be a good proxy for the expected stock return over the next period.

Second, equation (2) shows that in the absence of price bubbles, the log dividend-price ratio will be stationary if stock returns and dividend growth are stationary, conditions that seem quite plausible. In particular, if returns and dividend growth rates do not have time trends, then the log dividend-price ratio will not have a time trend either. (This model cannot be used to say what would happen if there were time trends in returns or dividend growth rates, because such trends would invalidate the linear approximation (1).) Third, however, persistent variation in returns or dividend growth rates can lead to persistent variation in the log dividend-price ratio even if that ratio is stationary.

The effect of persistence on predictive regressions has been highlighted by Stambaugh (1999). Stambaugh discusses the two-equation system,

$$r_{t-1} = \alpha + \beta x_t + u_{t+1} \tag{3}$$

$$x_{t-1} = \mu + \phi x_t + \eta_{t+1}, \tag{4}$$

where $x_t$ can be any persistent predictor variable but attention focuses on the level or log of the dividend-price ratio.

OLS estimates of equation (3) in twentieth-century U.S. data, with the log dividend-price ratio $x_t = d_t - p_t$ as the explanatory variable and the annualized stock return as the dependent variable, tend to deliver estimates in the range 0.1 to 0.2. An estimate of 0.04, the historical average level of the dividend-price ratio, would imply that around the average, a percentage point increase in the level of the dividend-price ratio increases the expected stock return by one percentage point. The OLS estimates imply a sensitivity of the return to the dividend-price ratio that is several times greater than this. They imply that when the dividend-price ratio is unusually high, it tends to return to normal through increases in prices that magnify the effect on stock returns. Campbell and Shiller (1998) emphasize this pattern in the historical data.

To understand Stambaugh's concern about persistence, define

$$\gamma = \frac{\sigma_{u\eta}}{\sigma_\eta^2}. \tag{5}$$

The coefficient $\gamma$ is the regression coefficient of return innovations on innovations to the predictor variable. In the case where the explanatory variable is the log dividend-price ratio, $\gamma$ is negative because rising stock prices tend to be associated with a falling dividend-price ratio. More precisely, dividend growth is only weakly correlated with and much less volatile than stock returns, so from equation (1) $\gamma$ is about $-\rho$, that is, slightly greater than $-1$.

Stambaugh points out that the bias in estimating the coefficient $\beta$ is $\gamma$ times the bias in estimating the persistence of the predictor variable, $\phi$:

$$\mathbf{E}[\hat{\beta} - \beta] = \gamma \mathbf{E}[\hat{\phi} - \phi]. \tag{6}$$

This is significant because it has been understood since the work of Kendall (1954) that there is downward bias in estimates of $\phi$ of about $-(1 + 3\phi)/T$, where $T$ is the sample size, primarily resulting from the fact that $x_t$ has an unknown mean that must be estimated. With a highly persistent predictor variable and $\gamma$ slightly greater than $-1$, the Stambaugh bias in $\hat{\beta}$ is almost $4/T$. With 50 years of data the bias is almost 0.08, substantial relative to the OLS estimates discussed above.

Recent responses to Stambaugh's critique have all used theory in one way or another. Lewellen (2004) first writes an expression for the bias conditional on the estimated persistence $\hat{\phi}$ and the true persistence $\phi$:

$$\mathrm{E}[\hat{\beta} - \beta \mid \hat{\phi}, \phi] = \gamma[\hat{\phi} - \phi]. \tag{7}$$

At first sight this expression does not seem particularly useful because we do not know the true persistence coefficient. However, Lewellen argues on the basis of theory that $\phi$ cannot be larger than one – the dividend-price ratio is not explosive – so the largest bias occurs when $\phi = 1$. He proposes the conservative approach of adjusting the estimated coefficient using this worst-case bias:

$$\hat{\beta}_{adj} = \hat{\beta} - \gamma(\hat{\phi} - 1). \tag{8}$$

In the data, the log dividend-price ratio appears highly persistent. That is, $\hat{\phi}$ is close to one; Lewellen reports a monthly estimate of 0.997 for the period 1946–2000, or about 0.965 on an annual basis. Lewellen's bias adjustment is therefore about 0.035, much smaller than Stambaugh's bias adjustment for a 50-year sample and somewhat smaller whenever the sample size is less than 114 years. Lewellen argues that stock returns are indeed predictable from the log dividend-price ratio, almost as much so as a naive researcher, unaware of Stambaugh's critique, might believe. Another way to express Lewellen's point is that data samples with spurious return predictability are typically samples in which the log dividend-price ratio appears to mean-revert more strongly than it truly does. In the historical data, the log dividend-price ratio has a root very close to unity – it barely seems to mean-revert at all – and thus we should not expect important spurious predictability in the historical data.

Cochrane (2007) responds to Stambaugh by directing attention to the inability of the log-dividend price ratio to forecast dividend growth. At first sight this response does not seem connected to Lewellen's, but in fact it is closely related. The Campbell-Shiller loglinearization (1) implies that $r_{t+1}$, $\Delta d_{t+1}$, $d_{t+1} - p_{t+1}$, and $d_t - p_t$ are deterministically linked. It follows that if we regress $r_{t+1}$, $\Delta d_{t+1}$, and $d_{t+1} - p_{t+1}$ onto $d_t - p_t$, the coefficients $\beta$, $\beta_d$, and $\phi$ are related by

$$\beta = 1 - \rho\phi + \beta_d, \tag{9}$$

where $\rho$ is the coefficient of loglinearization from equation (1).

If we have prior knowledge about $\phi$, then $\beta$ and $\beta_d$ are linked. For example, if $\rho = 0.96$ and we know that $\phi \leq 1$, then $\beta_d \leq \beta - 0.04$. If $\beta = 0$, then $\beta_d$ must be negative and less than $-0.04$. The fact that regression estimates of $\beta_d$ are close to zero is therefore indirect evidence that $\beta > 0$, in other words that stock returns are predictable – given our prior knowledge, based on theory, that the log dividend-price ratio is not explosive.

Another way to express Cochrane's point is that if the dividend-price ratio fails to predict stock returns, it will be explosive unless it predicts dividend growth. Since the dividend-price ratio cannot be explosive, the absence of predictable dividend growth strengthens the evidence for predictable returns.

Campbell and Yogo (2006) offer a third response to Stambaugh. They point out that if we knew persistence, we could reduce noise by adding the innovation to the predictor variable to the predictive regression, estimating

$$r_{t+1} = \alpha' + \beta x_t + \gamma(x_{t+1} - \phi x_t) + v_{t+1}. \tag{10}$$

The additional regressor, $(x_{t+1} - \phi x_t) = \eta_{t+1}$, is uncorrelated with the original regressor $x_t$ but correlated with the dependent variable $r_{t+1}$. Thus, the regression (10) still delivers a consistent estimate of the original predictive coefficient $\beta$, but it does so with increased precision because it controls for some of the noise in unexpected stock returns.

Of course, in practice we do not know the persistence coefficient $\phi$, but Campbell and Yogo argue that we can construct a confidence interval for it by inverting a unit root test. By doing this we 'de-noise' the return and get a more powerful test. The test delivers particularly strong evidence for predictability if we rule out a persistence coefficient $\phi > 1$ on prior grounds.

A way to understand Campbell and Yogo's results is to recall the challenge posed by the late 1990s. In that period, the dividend-price ratio was low, which led Campbell and Shiller (1998) to predict low stock returns based on a regression like (3). In fact, stock returns remained high until the early 2000s. These high returns were accompanied by falling dividend yields, despite the fact that the dividend yield was already below its historical mean. If we believe that the dividend yield was below its true mean and that it should be forecast to return to that mean rather than exploding away from it, then the late 1990s declines in the dividend-price ratio must have been unexpected. Unexpected declines in the dividend-price ratio are associated with unexpected high stock returns, accounting for the poor performance of the basic predictability regression in the late 1990s. The regression (10) corrects for this effect, limiting the negative influence of the late 1990s on the estimated predictive coefficient $\beta$.

The econometric issues discussed in this section have little effect on regressions that use nominal interest rates or yield spreads to predict excess stock returns. Although nominal interest rates are highly persistent, their innovations are not strongly correlated with innovations in stock returns, and thus the coefficient $\gamma$ is close to zero for these variables, implying only a trivial bias in OLS regression estimates. Even papers that are sceptical of stock return predictability from the dividend-price ratio, such as Ang and Bekaert (2007), emphasize the strength of the statistical evidence that interest rates predict stock returns. The challenge in this case is primarily a theoretical one: to understand the economic forces that cause common variation in nominal interest rates and the equity premium.

All the papers discussed above combine prior knowledge with classical statistical methods. It is possible, of course, to use finance theory in an explicit Bayesian manner. Several recent papers have done this, notably Pastor and Stambaugh (2007) and Wachter and Warusawitharana (2007). Consistent with the results reported here, these papers find that tight priors on the persistence of the predictor variable tend to deliver stronger evidence for predictability of stock returns.

## 3. Steady-state return prediction

The papers discussed in the previous section address the question of whether the equity premium varies with market valuations, or whether it is constant. Even if one believes that the equity premium is time varying, however, there remains the important question of how best to estimate it at each point in time. Given the noise in stock returns, equity premium models with multiple free coefficients are hard to estimate and may fail out of sample because of errors in estimating the coefficients. Indeed, Goyal and Welch (2007) argue that almost all the regression models proposed in the recent literature fail to beat the historical sample mean when predicting excess stock returns out of sample.

In response to Goyal and Welch, Campbell and Thompson (2007) propose to use steady-state valuation models to estimate the equity premium. Such models tightly restrict the way in which historical data are used to predict future returns, and Campbell and Thompson find that they work well out of sample. Fama and French (2002) and Pastor, Sinha, and Swaminathan (2007) also use this approach to analyse the equity premium. The approach is analogous to the familiar procedure of forecasting the return on a bond, using its yield rather than its historical average return.

The classic steady-state model is the Gordon growth model, named after Canadian economist Myron Gordon. The model describes the level of the dividend-price ratio in a steady state with a constant discount rate and growth rate. Using upper-case letters to denote levels of variables, the Gordon growth model can be written as

$$\frac{D}{P} = R - G. \tag{11}$$

This formula can be used directly with historical dividend growth rates, but it can also be rewritten in several ways that suggest alternative empirical strategies for forecasting stock returns. First, one can substitute out growth by using the steady-state relation between growth and accounting return on equity,

$$G = \left(1 - \frac{D}{E}\right) ROE, \tag{12}$$

where $D/E$ is the payout ratio, to obtain a growth-adjusted return forecast

$$\hat{R}_{DP} = \frac{D}{P} + \left(1 - \frac{D}{E}\right) ROE. \tag{13}$$

This return forecast is linear in $D/P$, with a slope coefficient of one and an intercept that is determined by the reinvestment rate and profitability. Importantly, neither the slope coefficient nor the intercept need to be estimated from noisy historical stock returns.

Second, one can restate the model in terms of the earnings-price ratio by using $D/P = (D/E)(E/P)$ to obtain

$$\hat{R}_{EP} = \left(\frac{D}{E}\right)\frac{E}{P} + \left(1 - \frac{D}{E}\right) ROE, \tag{14}$$

a payout-ratio-weighted average of the earnings-price ratio and the accounting return on equity. When return on equity equals the expected return, as might be the case in long-run equilibrium, then this implies that $\hat{R}_{EP} = E/P$.

Finally, one can rewrite the model in terms of the book-market ratio. Since $E/P = (B/M)ROE$,

$$\hat{R}_{BM} = ROE\left[1 + \frac{D}{E}\left(\frac{B}{M} - 1\right)\right]. \tag{15}$$

To use these formulas in practice, one must decide how to combine historical and contemporaneous data on the right-hand-side variables. Campbell and Thompson (2007) follow Fama and French (2002) by using historical average data on payouts and profitability, but differ from them by using current rather than historical average data on valuation ratios to obtain a return forecast conditional on the market's current valuation level. This procedure assumes that movements in valuation ratios, relative to historical cash flows, are explained by permanent changes in expected returns, so that each percentage point increase in the level of the dividend-price ratio generates a percentage point increase in the return forecast. It is a compromise between the view that valuation ratios are driven by changing forecasts of profitability, in which case the implied movements in returns would be smaller, and the view that valuation ratios are driven by *temporary* changes in discount rates, in which case the implied return movements would be larger, as discussed in the previous section.

Campbell and Thompson evaluate the out-of-sample performance of these models and several other variants over the period 1927–2005 and subsamples with breakpoints at 1956 and 1980. They find that steady-state valuation models typically perform better when more theoretical restrictions are imposed, and that they almost always outperform the historical mean return as a predictor of future returns. Dividend-based and earnings-based models, equations (13) and (14), generally appear to be more successful than the book-market model (15). In the next section I illustrate this approach using a model that averages both the dividend-price ratio and the recent history of earnings to generate a return forecast that is a blend of those from (13) and (14).

### *3.1. The Gordon model with a random walk dividend-price ratio*

It may at first sight appear strange that steady-state valuation models based on the Gordon growth model perform well, given that they assume constant valuation ratios, while in the data valuation ratios vary in a highly persistent manner. It turns out, however, that a variant of the Gordon growth model can be derived using the assumption that the log dividend-price ratio follows a random walk. Under this assumption the Campbell-Shiller loglinear model, used in the previous section, breaks down because the dividend-price ratio has no fixed mean around which to take a loglinear approximation. However, in this case a suitable version of the original Gordon growth model is available to take the place of the Campbell-Shiller model.

To show this I assume, as in the Gordon growth model, that the dividend is known one period in advance. Then we can write

$$\frac{D_{t+1}}{P_t} = \exp(x_t), \tag{16}$$

where $x_t$ now denotes the log dividend-price ratio using a forward or indicated dividend rather than a historical dividend. I assume that $x_t$ follows a random walk:

$$x_t = x_{t-1} + \varepsilon_t. \tag{17}$$

Since the dividend growth rate is known one period in advance, I can write

$$\frac{D_{t+1}}{D_t} = 1 + G_t = \exp(g_t). \tag{18}$$

Finally, I assume that $x_{t-1}$ and $g_{t+1}$ are conditionally normal given time $t$ information.

The definition of the stock return implies that

$$\begin{aligned} 1 + R_{t-1} &= \frac{P_{t-1} + D_{t-1}}{P_t} = \frac{D_{t+1}}{P_t} + \frac{D_{t+2}}{D_{t+1}} \frac{D_{t-1}}{P_t} \left( \frac{D_{t-2}}{P_{t-1}} \right)^{-1} \\ &= \exp(x_t)[1 + \exp(g_{t+1} - x_{t-1})]. \end{aligned} \tag{19}$$

The conditionally expected stock return can be calculated using the formula for the conditional expectation of lognormally distributed random variables and the martingale property that $\mathrm{E}_t\, x_{t-1} = x_t$:

$$
\begin{aligned}
E_t(1+R_{t+1}) &= \exp(x_t)[1+\mathrm{E}_t \exp(g_{t+1}-x_{t+1})] \\
&= \exp(x_t)\left[1+\exp\left(\mathrm{E}_t g_{t+1}-x_t+\sigma_g^2/2+\sigma_x^2/2-\sigma_{gx}\right)\right] \\
&= \frac{D_{t-1}}{P_t}+\exp(\mathrm{E}_t g_{t+1})\exp(\mathrm{Var}_t(p_{t+1}-p_t)/2).
\end{aligned} \tag{20}
$$

Finally, the right-hand side of (20) can be approximated using the facts that for small $y$, $\exp(y) \approx 1+y$, and that unexpected log stock returns are approximately equal to unexpected changes in log stock prices:

$$
E_t(1+R_{t+1}) \approx \frac{D_{t-1}}{P_t}+\exp(\mathrm{E}_t g_{t+1})+\frac{1}{2}\mathrm{Var}_t(r_{t+1}). \tag{21}
$$

This equation expresses the expected stock return as the level of the dividend yield, plus geometric average dividend growth, plus one-half the variance of stock returns. In the original Gordon model, $\sigma_x^2 = 0$, so the variance of stock returns equals the variance of dividend growth. Since arithmetic average dividend growth equals geometric average dividend growth plus one-half the variance of dividend growth, in this case we get the original Gordon formula that the arithmetic average stock return equals dividend yield plus arithmetic average dividend growth.

If one subtracts half the variance of stock returns from each side of (20), one finds that the geometric average stock return equals the level of the dividend-price ratio plus the geometric average of dividend growth. Under the assumptions of the original Gordon model, the geometric implementation of the model is equivalent to an arithmetic implementation because stock returns and dividend growth have the same variance, so their geometric and arithmetic averages differ by the same amount. In the data, however, returns are much more volatile, so the geometric implementation and the arithmetic implementation are different. The analysis here shows that the geometric implementation is correct. Interestingly, this is exactly the way in which the model is used by Siegel (1994).

## 4. What is the equity premium today?

I now use a version of the above methodology, starting from equation (14), to estimate the equity premium. Following the previous discussion, I first estimate the conditional geometric average stock return, then subtract the real interest rate to get an equity premium number, and finally discuss the adjustment that is needed to convert from a geometric average to an arithmetic average equity premium. I look at data for the world as a whole (measured using the Morgan Stanley Capital International all-world index), and also for the U.S. and Canada, over the period from 1982 through the end of March 2007.

Figure 1 shows that for all three indices smoothed earnings-price ratios, with earnings smoothed over three years to eliminate cyclical noise, have fallen

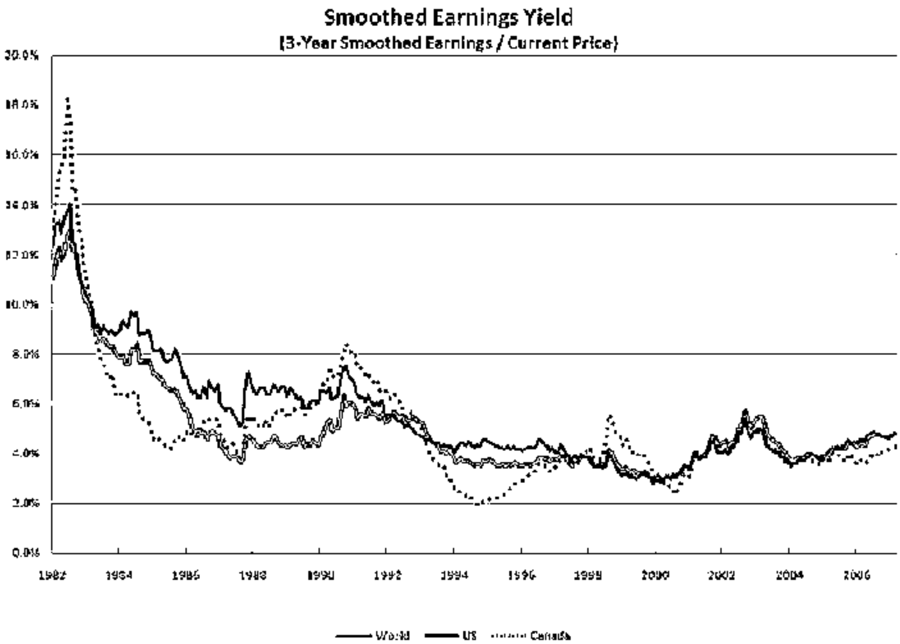


FIGURE 1 Three-year smoothed earnings-price ratios in the world, the U.S., and Canada

dramatically since the early 1980s and have been in the 3% to 5% range for the last ten years. During the same period, however, figure 2 shows that profitability has increased from a long-run historical average of around 6% to much higher values around 10%. Meanwhile, payout ratios have fluctuated widely around an average of about 50%.

In constructing a return forecast, it is desirable to combine historical earnings with some forward-looking measure of earnings. One possibility is to use analysts' earnings forecasts (Pastor, Sinha, and Swaminathan 2007); another is to use dividends. I average historical earnings, smoothed over three years, and the current dividend, divided by the payout rate, to construct a forward-looking measure of permanent earnings that can be used in equation (14).

When I put these numbers together, an earnings-based estimate of the real return on U.S. equities, assuming constant 6% real profitability and a 50% payout rate, was about 9% in the early 1980s and fell to just above 4% in the year 2000. Since then it has increased to slightly over 5%. This estimate assumes that profitability and payouts are best forecast to be constant; alternatively, if one uses the three-year moving average of profitability illustrated in figure 2, and a similar three-year moving average of the payout ratio, the current real return estimate increases by almost 4% to 9%, reflecting the high recent profitability and low payout ratios of U.S. corporations. At the world level, the current real return number is comparable to the U.S. number if a fixed profitability estimate

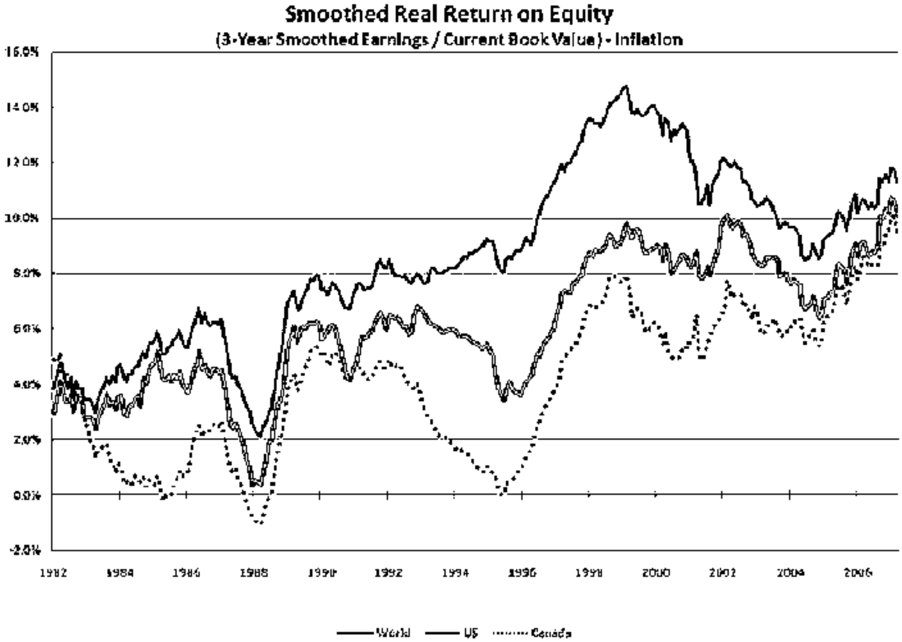


FIGURE 2 Three-year smoothed profitability in the world, the U.S., and Canada

is used, but the adjustment for recent profitability and payouts is much smaller, only slightly above 2%. The Canadian real return number is also very similar to that in the U.S. on the basis of fixed profitability, but lower Canadian profitability and higher payouts in the last few years imply that the use of recent data increases the estimated real return by less than 2%.

To convert these numbers into estimates of the equity premium, one needs to subtract a safe real interest rate. Figure 3 plots real yields on inflation-indexed bonds in three large markets, the U.K., the U.S., and Canada. The figure shows that the average real yield on inflation-indexed bonds across the three countries was about 3.5% in the 1990s but fell below 2% in the early 2000s. By the end of March 2007, it had recovered to just over 2%.

The implied current equity premium, assuming constant profitability and payouts, is just over 3%: 3.3% for the world as a whole, 3.2% for the U.S., and 3.1% for Canada. If instead one uses recent profitability and payouts, the current equity premium is 5.7% for the world as a whole, a startling 6.9% for the U.S., and 5.0% for Canada. Figures 4, 5, and 6 illustrate the history of the equity premium in the world, the U.S., and Canada under these two alternative assumptions.

Obviously a key question is whether the high profitability of global, and particularly U.S., corporations can be expected to continue. On the one hand, globalization has increased the supply of labour relative to capital, reducing wage pressure and increasing profitability; on the other hand, profitability has been increased

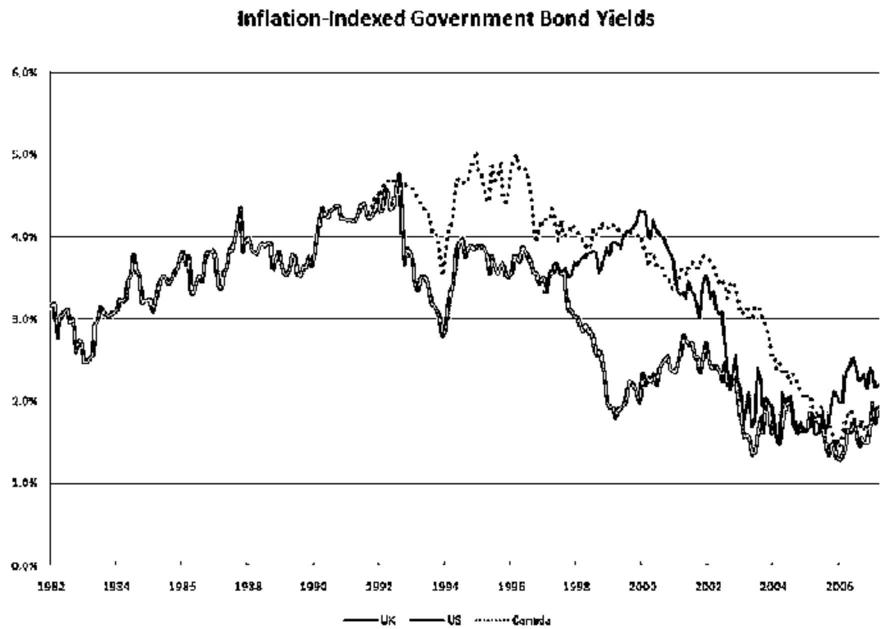


FIGURE 3 Long-term real interest rates in the U.K., the U.S., and Canada

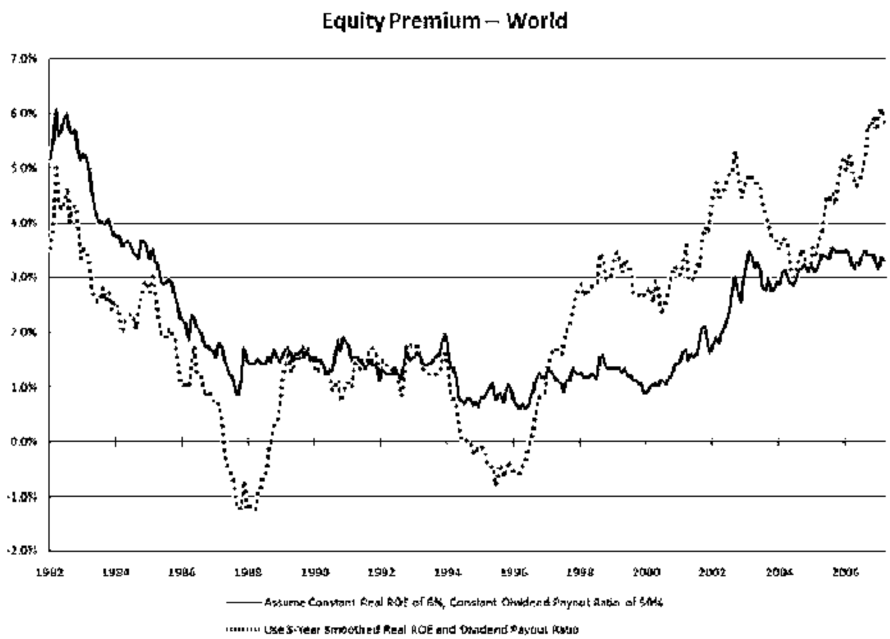


FIGURE 4 The world equity premium since 1982

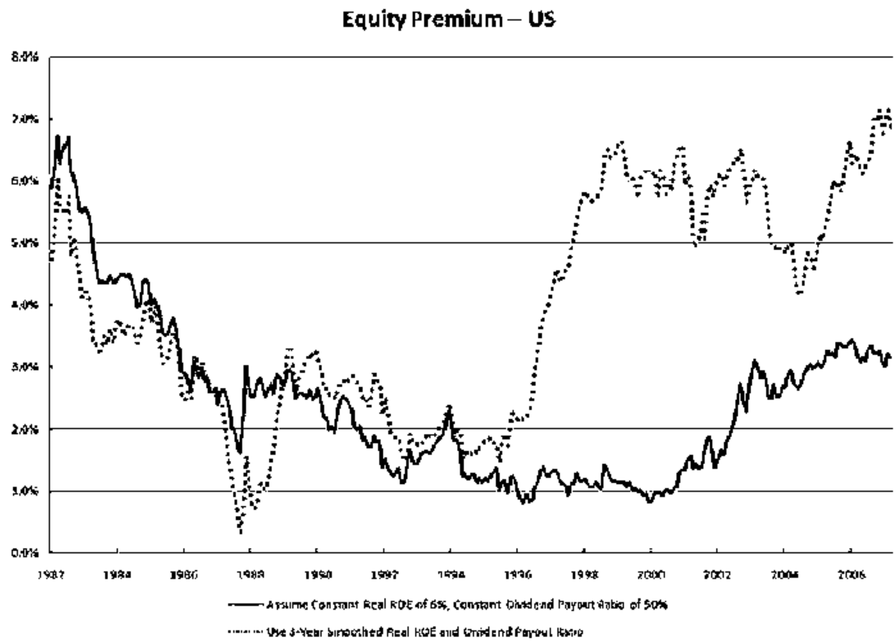


FIGURE 5 The U.S. equity premium since 1982

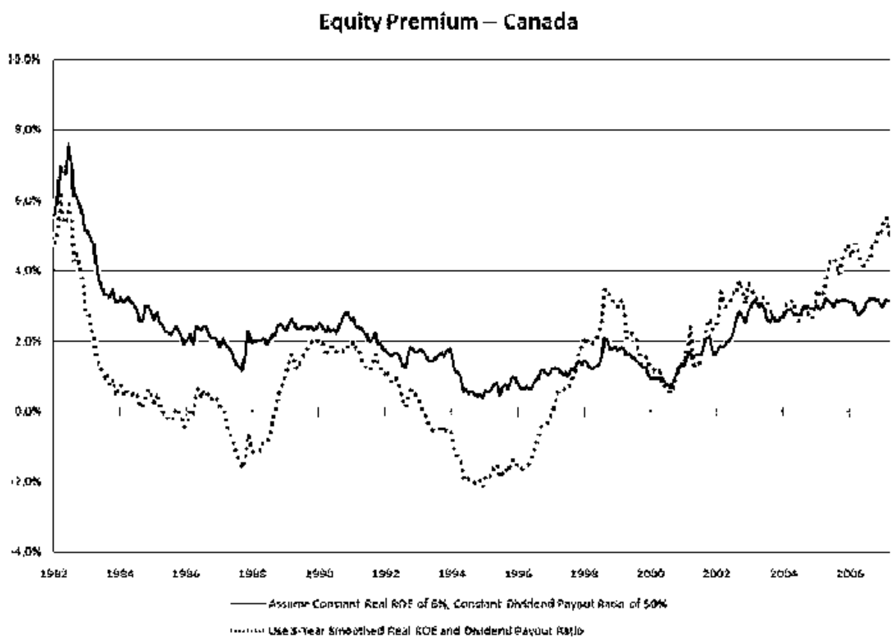


FIGURE 6 The Canadian equity premium since 1982

by favourable business cycle and political conditions that may not persist. Historically, profitability has shown temporary fluctuations and low payout rates (high reinvestment rates) have predicted declining profitability. Also, equity premium estimates based on current profitability and payout rates have been highly volatile, even turning negative on occasion. For both these reasons it seems wise to place considerably more weight on long-term averages than on recent data. If one puts a weight of 0.75 on the long-term average, with 0.25 on the recent data, the implied equity premium at the end of March 2007 is in the range 3.6% to 4.1%: 3.9% in the world as a whole, 4.1% in the U.S., and 3.6% in Canada. This number is a geometric average equity premium; for an arithmetic average, one should add one-half the variance of stock returns, or almost 1.3% if stock returns have a conditional standard deviation of 16%. The resulting arithmetic equity premium numbers are in the range 4.9% to 5.4%. Note that the equity premium is this high in large part because the safe real interest rate has declined over the past decade, as illustrated in figure 3.

These numbers are lower than historical average excess stock returns reported by Dimson, Marsh, and Staunton (2006). Using data for the period 1900–2005, Dimson, Marsh, and Staunton report geometric average equity premia of 4.7% for the world as a whole, 5.5% for the U.S., and 4.5% for Canada. The difference reflects two facts. First, historical average returns have been driven up by declining valuation ratios; this effect cannot be expected to continue in the future because valuation ratios should not have trends, a point emphasized by Fama and French (2002). Second, historical average returns were obtained by investors who paid lower stock prices and thus benefited from higher dividend-price ratios.

It is interesting to note that chief financial officers of major corporations, surveyed by Graham and Harvey (2007), have modest expectations of the equity premium, which implies that they do not expect recent profitability to continue. Their median estimate of the geometric average U.S. equity premium at the end of November 2006 was 3.4%, much closer to the constant-profitability number reported here than to the recent-profitability number and far below the historical average equity premium.

## 5. Return prediction with cross-sectional variables

Finance theory can also be used to predict excess stock returns using information in the cross-section of stock prices. This is valuable both to corroborate the predictions from aggregate valuation ratios and possibly as a way to pick up higher-frequency components of the equity premium that may be missed by a steady-state approach.

Polk, Thompson, and Vuolteenaho (2006) argue that if the Capital Asset Pricing Model (CAPM) is true, then a high equity premium implies low prices for stocks that have high betas with the aggregate market index. That is, high-beta stocks should be value stocks with low ratios of market prices to accounting measures of fundamental value. Reversing the argument, value stocks should tend to have high betas. This was true in the mid-twentieth-century, roughly from the 1930s through the 1950s, but in recent decades growth stocks have had higher betas than value stocks (Franzoni 2006). Polk, Thompson, and Vuolteenaho argue that this change in cross-sectional stock pricing reflects a decline in the equity premium. They construct a predictor of the aggregate market return, based on the relative pricing of high- and low-beta stocks, and show that it correlates well with the smoothed earnings-price ratio except in the early 1980s when inflation may have distorted the relationship.

It is possible to push this idea even further, exploiting the fact that the CAPM may not fully describe the cross-section of stock returns when returns are predictable in the time series. Merton (1973) developed an intertemporal CAPM (ICAPM) that showed that in the presence of time-varying expected returns, long-lived investors care not only about shocks to their wealth but also about shocks to the expected return on wealth. Intuitively, they value wealth not for its own sake but for the consumption stream it can provide; thus, they want to hedge against declines in the rate of return just as much as against declines in market value. Campbell (1993) implemented this idea using a vector autoregression (VAR) to break market movements into permanent movements driven by news about cash flows and temporary movements driven by news about discount rates. Long-lived investors are more concerned about the former than about the latter. Thus, stocks that covary with cash-flow news should have higher average returns than stocks that covary with discount-rate news, when betas with the overall market return are controlled for.

One of the main deviations from the CAPM in recent decades has been the value effect, the high average returns that value stocks have delivered despite their low market betas. If the ICAPM is to explain the value effect, it must be that value stocks covary with cash-flow news while growth stocks covary with discount-rate news. This implies that a moving average of past excess returns on growth stocks should be a good predictor of aggregate stock returns.

The value spread, the relative valuation of value and growth stocks (normally measured as the difference between the log book-market ratios of these two types of stocks) is one possible summary of past excess returns on growth stocks. Eleswarapu and Reinganum (2004) find that the value spread for small stocks predicts the aggregate market return, and Campbell and Vuolteenaho (2004) use the same variable in a VAR model to estimate and test the ICAPM. They find that the ICAPM explains the average returns of value and growth stocks much better than does the standard CAPM. Cohen, Polk, and Vuolteenaho (2006) and Campbell, Polk, and Vuolteenaho (2007) explore the robustness of these results, using both VAR-based and direct measures of cash-flow and discount-rate news. Empirically, the effect of including the small-stock value spread in a model of the equity premium is to lower the estimated equity premium at the turn of the millennium, when growth stocks were abnormally expensive relative to value stocks, and to increase it in 2006 and early 2007, when growth stocks were abnormally cheap.

All this work relies on theoretically motivated, but not fully restricted, time-series models of the aggregate market return. A natural next step is to use the theoretical restrictions of the ICAPM to jointly estimate a time-series model of the aggregate market return and a cross-sectional model of average stock returns. Campbell (1996) was an early implementation of this approach, but that paper did not find systematic deviations from the CAPM because it did not use the information in the relative prices of growth and value stocks. Recent research suggests that with the proper information variables and test assets, cross-sectional information can play an important role in a jointly estimated model of the equity premium.

## 6. Conclusion

In this paper I have tried to illustrate the usefulness of finance theory for statistical analysis of stock returns, in particular for estimation of the equity premium. The literature on this topic is vast, and inevitably I have neglected some important aspects. Five omissions deserve special mention.

First, I have not reviewed the simple but important point that excess stock returns should be difficult to predict, because highly predictable excess returns would imply extremely large profits for market-timing investors. Campbell and Thompson (2007) explore the mapping from $R^2$ statistics in predictive regressions to profits and welfare gains for market timers. The basic lesson is that investors should be suspicious of predictive regressions with high $R^2$ statistics, asking the old question, 'If you're so smart, why aren't you rich?'

Second, I have confined attention to short-term predictive regressions and have not considered direct forecasts of long-horizon returns. It has been known since Fama and French (1988) that long-horizon regressions often have higher $R^2$ statistics than short-horizon regressions, but their statistical properties are controversial. Campbell (2001) and Cochrane (2007) argue that in certain circumstances, long-horizon regressions can have superior power to detect predictability when in fact it exists.

Third, I have not discussed recent work that uses finance theory to infer the equity premium from the actions of market participants. Lettau and Ludvigson (2001), for example, argue that the level of consumption in relation to aggregate financial wealth and labour income reveals consumers' expectations of future stock returns. In a similar spirit Baker and Wurgler (2000) use the financing decisions of corporations to infer corporate managers' beliefs about expected stock returns.

Fourth, I have presented estimates of the equity premium without discussing the uncertainty of these estimates. I have suggested that finance theory can reduce our uncertainty about the equity premium, but a more formal Bayesian analysis would be needed to quantify this effect.

Finally, I have not attempted to review the important body of empirical work on the estimation of stock market risk. Mechanically, the volatility of stock returns determines the wedge between geometric and arithmetic average stock returns. Economically, both risk and return matter to investors, and it is plausible that changing risk is one factor that drives the changing equity premium. Merton (1980), Campbell (1987), French, Schwert, and Stambaugh (1987), Harvey (1989), and Glosten, Jagannathan, and Runkle (1993) are a few of the earlier papers that explore this relation. Recent contributions by Ghysels, Santa-Clara, and Valkanov (2005) and Pastor, Sinha, and Swaminathan (2007) find that the equity premium does covary positively with estimated risk, but that this effect does not explain the predictability of stock returns from valuation ratios or interest rates.

Despite the size and complexity of the literature on the equity premium, it has a simple unifying theme. Campbell, Lo, and MacKinlay (1997) argue that 'what distinguishes financial economics is the central role that uncertainty plays in both financial theory and its empirical implementation.' Theory tells us why stock returns are so hard to predict. But it also holds out the promise of better prediction than we can hope to achieve by purely statistical forecasting methods.

## References

Ang, Andrew, and Geert Bekaert (2007) 'Stock return predictability: is it there?' *Review of Financial Studies* 20, 651–707

Baker, Malcolm, and Jeffrey Wurgler (2000) 'The equity share in new issues and aggregate stock returns,' *Journal of Finance* 55, 2219–57

Boudoukh, Jacob, Roni Michaely, Matthew Richardson, and Michael Roberts (2007) 'On the importance of measuring payout yield: Implications for empirical asset pricing,' *Journal of Finance* 62, 877–916

Butler, Alexander W., Gustavo Grullon, and James P. Weston (2005) 'Can managers forecast aggregate market returns?' *Journal of Finance* 60, 963–86

Campbell, John Y. (1987) 'Stock returns and the term structure,' *Journal of Financial Economics* 18, 373–99

— (1993) 'Intertemporal asset pricing without consumption data,' *American Economic Review* 83, 487–512

(1996) 'Understanding risk and return,' *Journal of Political Economy* 104, 298–345

(2001) 'Why long horizons? A study of power against persistent alternatives,' *Journal of Empirical Finance* 8, 459–91

Campbell, John Y., and John H. Cochrane (1999) 'By force of habit: A consumption-based explanation of aggregate stock market behavior,' *Journal of Political Economy* 107, 205–51

Campbell, John Y., and Robert J. Shiller (1988a) 'The dividend-price ratio and expectations of future dividends and discount factors,' *Review of Financial Studies* 1, 195–228

(1988b), 'Stock prices, earnings, and expected dividends,' *Journal of Finance* 43, 661–76

— (1998). 'Valuation ratios and the long-run stock market outlook,' *Journal of Portfolio Management* 24(2), 11–26

Campbell, John Y., and Samuel B. Thompson (2007) 'Predicting excess stock returns out of sample: Can anything beat the historical average?' *Review of Financial Studies*, forthcoming

Campbell, John Y., and Tuomo O. Vuolteenaho (2004) 'Bad beta, good beta,' *American Economic Review* 94, 1249–75

Campbell, John Y., and Motohiro Yogo (2006) 'Efficient tests of stock return predictability,' *Journal of Financial Economics* 81, 27–60

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay (1997) *The Econometrics of Financial Markets* (Princeton, NJ: Princeton University Press)

Campbell, John Y., Christopher Polk, and Tuomo O. Vuolteenaho (2007) 'Growth or glamour? Fundamentals and systematic risk in stock returns,' working paper, Harvard University and London School of Economics

Cavanagh, Christopher L., Graham Elliott, and James H. Stock (1995) 'Inference in models with nearly integrated regressors,' *Econometric Theory* 11, 1131–47

Cochrane, John H. (2007) 'The dog that did not bark: a defense of return predictability,' *Review of Financial Studies*, forthcoming

Cohen, Randolph B., Christopher Polk, and Tuomo O. Vuolteenaho (2006) 'The price is (almost) right,' working paper, Harvard University and Northwestern University

Dimson, Elroy, Paul Marsh, and Michael Staunton (2006) 'The worldwide equity premium: a smaller puzzle,' working paper, London Business School

Eleswarapu, Venkat R. and Marc R. Reinganum (2004) 'The predictability of aggregate stock market returns: evidence based on glamour stocks,' *Journal of Business* 77, 275–94

Fama, Eugene F., and Kenneth R. French (1988) 'Dividend yields and expected stock returns,' *Journal of Financial Economics* 22, 3–25

(1989) 'Business conditions and expected returns on stocks and bonds,' *Journal of Financial Economics* 25, 23–49

— (2002) 'The equity premium,' *Journal of Finance* 57, 637–59

Fama, Eugene F., and G. William Schwert (1977) 'Asset returns and inflation,' *Journal of Financial Economics* 5, 115–46

Ferson, Wayne E., Sergei Sarkissian, and Timothy T. Simin (2003) 'Spurious regressions in financial economics?' *Journal of Finance* 58, 1393–413

Foster, F. Douglas, Tom Smith, and Robert E. Whaley (1997) 'Assessing goodness-of-fit of asset pricing models: the distribution of the maximal $R^2$,' *Journal of Finance* 52, 591–607

Franzoni, Francesco (2006) 'Where is beta going? The riskiness of value and small stocks,' working paper, HEC School of Management

French, Kenneth R., G. William Schwert, and Robert F. Stambaugh (1987) 'Expected stock returns and volatility,' *Journal of Financial Economics* 19, 3–29

Ghysels, Eric, Pedro Santa-Clara, and Rossen Valkanov (2005) 'There is a risk-return tradeoff after all,' *Journal of Financial Economics* 76, 509–48

Glosten, Lawrence R., Ravi Jagannathan, and David E. Runkle (1993) 'On the relation between the expected value and the volatility of the nominal excess return on stocks,' *Journal of Finance* 48, 1779–801

Goyal, Amit, and Ivo Welch (2003) 'Predicting the equity premium with dividend ratios,' *Management Science* 49, 639–54

(2007) 'A comprehensive look at the empirical performance of equity premium prediction,' *Review of Financial Studies*, forthcoming

Graham, Benjamin, and David L. Dodd (1934) *Security Analysis*, 1st ed. (New York: McGraw-Hill)

Graham, John, and Campbell Harvey (2007) 'The equity risk premium in January 2007: evidence from the global CFO outlook survey,' working paper, Duke University

Harvey, Campbell (1989) 'Time-varying conditional covariances in tests of asset pricing models,' *Journal of Financial Economics* 24, 289–317

Hodrick, Robert J. (1992) 'Dividend yields and expected stock returns: alternative procedures for inference and measurement,' *Review of Financial Studies* 5, 257–86

Jansson, Michael, and Marcelo J. Moreira (2006) 'Optimal inference in regression models with nearly integrated regressors,' *Econometrica* 74, 681–715

Keim, Donald B., and Robert F. Stambaugh (1986) 'Predicting returns in the stock and bond markets,' *Journal of Financial Economics* 17, 357–90

Kendall, Maurice G. (1954) 'Note on bias in the estimation of autocorrelation,' *Biometrika* 41, 403–4

Kothari, S.P., and Jay Shanken (1997) 'Book-to-market, dividend yield, and expected market returns: a time-series analysis,' *Journal of Financial Economics* 44, 169–203

Lamont, Owen (1998) 'Earnings and expected returns,' *Journal of Finance* 53, 1563–87

Lettau, Martin, and Sydney Ludvigson (2001) 'Consumption, aggregate wealth, and expected stock returns,' *Journal of Finance* 56, 815–49

Lewellen, Jonathan (2004) 'Predicting returns with financial ratios,' *Journal of Financial Economics* 74, 209–35

Merton, Robert C. (1973) 'An intertemporal capital asset pricing model,' *Econometrica* 41, 867–87

— (1980) 'On estimating the expected return on the market: an exploratory investigation,' *Journal of Financial Economics* 8, 323–61

Nelson, Charles, and M. Kim (1993) 'Predictable stock returns: the role of small sample bias,' *Journal of Finance* 48, 641–61

Pastor, Lubos, and Robert F. Stambaugh (2007) 'Predictive systems: living with imperfect predictors,' NBER Working Paper 12814

Pastor, Lubos, Meenakshi Sinha, and Bhaskaran Swaminathan (2007) 'Estimating the intertemporal risk-return tradeoff using the implied cost of capital,' *Journal of Finance*, forthcoming

Polk, Christopher, Samuel Thompson, and Tuomo Vuolteenaho (2006) 'Cross-sectional forecasts of the equity premium,' *Journal of Financial Economics* 81, 101–41

Pontiff, Jeffrey, and Lawrence D. Schall (1998) 'Book-to-market ratios as predictors of market returns,' *Journal of Financial Economics* 49, 141–60

Rozeff, Michael S. (1984) 'Dividend yields are equity risk premiums,' *Journal of Portfolio Management* 11(1), 68–75

Siegel, Jeremy (1994) *Stocks for the Long Run* (New York: McGraw-Hill)

Stambaugh, Robert F. (1999) 'Predictive regressions,' *Journal of Financial Economics* 54, 375–421

Torous, Walter, Rossen Valkanov, and Shu Yan (2004) 'On predicting stock returns with nearly integrated explanatory variables,' *Journal of Business* 77, 937–66

Wachter, Jessica A., and Missaka Warusawitharana (2007) 'Predictable returns and asset allocation: should a skeptical investor time the market?' NBER Working Paper No.13165

# Estimating the Real Rate of Return on Stocks Over the Long Term

*Papers by*

*John Y. Campbell*
*Peter A. Diamond*
*John B. Shoven*

*Presented to the*
***Social Security Advisory Board***

*August 2001*

**Social Security Advisory Board**

An independent, bipartisan Board created by Congress and appointed by the President and the Congress to advise the President, the Congress, and the Commissioner of Social Security on matters related to the Social Security and Supplemental Security Income programs.

# TABLE OF CONTENTS

# INTRODUCTION

In recent years there have been a variety of proposals that would change the current Social Security system to include some form of investment of funds in private equities. These proposals include allowing or requiring individuals to use a portion of the payroll tax to fund individual investment accounts, either as part of the Social Security system or as an addition to it. They also include proposals to require the government to invest a portion of the Social Security Trust Funds in equities.

A key element in evaluating these proposals is the rate of return that can be expected on such investments. The members of the 1994-1996 Advisory Council on Social Security agreed to use a real annual rate of 7 percent (the average for the period 1900-1995) to compare the three plans put forward by the Council. The Office of the Chief Actuary (OCACT) of the Social Security Administration has continued to use 7 percent to evaluate proposals for investment in stocks. However, there is a question as to whether the historical rate for the last century should be used to make long-term projections over the coming decades or whether an alternative rate or range of rates is more appropriate.

This document includes papers by three distinguished economists that examine this important question, including the issue of how to reflect the higher risk inherent in stock investment relative to investment in U.S. Treasury securities. The papers are by John Campbell, Otto Eckstein Professor of Applied Economics at Harvard University; Peter Diamond, Institute Professor at the Massachusetts Institute of Technology; and John Shoven, Charles Schwab Professor of Economics at Stanford University. The Board is publishing them in order to make them available to policy makers and members of the public who are interested in the issue of how to ensure the long-term solvency of the Social Security system.

The papers (which have been updated for purposes of this document) were the basis for a discussion sponsored by the Social Security Advisory Board on May 31, 2001. The purpose of the discussion was to enable individuals from OCACT who have the responsibility of estimating the effects of changes in the Social Security system to hear a range of views on the likely real yields on equities over the long term. Participants in the discussion from OCACT included Stephen Goss, Chief Actuary; Alice Wade, Deputy Chief Actuary; Patrick Skirvin, Lead Economist; and Anthony Cheng, Economist.

Participants also included three other distinguished economists who were on the 1999 Technical Panel on Assumptions and Methods: Eugene Steuerle, Senior Fellow, The Urban Institute; Deborah Lucas, Professor of Finance, Northwestern University and currently Chief Economist, Congressional Budget Office; and Andrew Samwick, Assistant Professor of Economics, Dartmouth College. The 1999 Technical Panel, which was sponsored by the Advisory Board, was charged with reviewing the assumptions and methods used in the long-term projections of the Social Security Trust Funds. The Panel also examined the question of how to evaluate the returns and risks involved in stock market investments. The Panel's report was published by the Board in November 1999 and is available on the Board's Web site (www.ssab.gov).

# Forecasting U.S. Equity Returns in the 21st Century

**John Y. Campbell, Professor of Economics**
**Harvard University**
**July 2001**

What returns should investors expect the U.S. stock market to deliver on average during the next century? Does the experience of the last century provide a reliable guide to the future? In this short note I first discuss alternative methodologies for forecasting average future equity returns, then discuss current market conditions, and finally draw conclusions for long-term return forecasts. Throughout I work in real, that is inflation-adjusted, terms.

## I. Methods for Forecasting Returns

### 1. Average past returns

Perhaps the simplest way to forecast future returns is to use some average of past returns. Very naturally, this method has been favored by many investors and analysts. However there are several difficulties with it.

**a)** ***Geometric average or arithmetic average?*** The geometric average return is the cumulative past return on U.S. equities, annualized. Siegel (1998) studies long-term historical data on value-weighted U.S. share indexes. He reports a geometric average of 7.0% over two different sample periods, 1802-1997 and 1871-1997. The arithmetic average return is the average of one-year past returns on U.S. equities. It is considerably higher than the geometric average return, 8.5% over 1802-1997 and 8.7% over 1871-1997.[1]

When returns are serially uncorrelated, the arithmetic average represents the best forecast of future return in any randomly selected future year. For long holding periods, the best forecast is the arithmetic average compounded up appropriately. If one is making a 75-year forecast, for example, one should forecast a cumulative return of $1.085^{75}$ based on 1802-1997 data.

When returns are negatively serially correlated, however, the arithmetic average is not necessarily superior as a forecast of long-term future returns. To understand this, consider an extreme example in which prices alternate deterministically between 100 and 150. The return is 50% when prices rise, and -33% when prices fall. Over any even number of periods, the geometric average return is zero, but the arithmetic average return is 8.5%. In this case the arithmetic average return is misleading because it fails to take account of the fact that high returns always multiply a low initial price of 100, while low returns always multiply a high initial price of

[1] When returns are lognormally distributed, the difference between the two averages is approximately one-half the variance of returns. Since stock returns have an annual standard deviation of about 18% over these long periods, the predicted difference is $0.18^2/2=0.016$ or 1.6%. This closely matches the difference in the data.

150. The geometric average is a better indication of long-term future prospects in this example.[2]

This point is not just a theoretical curiosity, because in the historical data summarized by Siegel, there is strong evidence that the stock market is mean-reverting. That is, periods of high returns tend to be followed by periods of lower returns. This suggests that the arithmetic average return probably overstates expected future returns over long periods.

**b) *Returns are very noisy.*** The randomness in stock returns is extreme. With an annual standard deviation of real return of 18%, and 100 years of past data, a single year's stock return that is only one standard deviation above average increases the average return by 18 basis points. A lucky year that is two standard deviations above average increases the average return by 36 basis points. Even when a century or more of past data is used, forecasts based on historical average returns are likely to change substantially from one year to the next.

**c) *Realized returns rise when expected returns fall.*** To the extent that expected future equity returns are not constant, but change over time, they can have perverse effects on realized returns. Suppose for example that investors become more risk-tolerant and reduce the future return that they demand from equities. If expected future cash flows are unchanged, this drives up prices and realized returns. Thus an estimate of future returns based on average past realized returns will tend to increase just as expected future returns are declining.

Something like this probably occurred in the late 1990's. A single good year can have a major effect on historical average returns, and several successive good years have an even larger effect. But it would be a mistake to react to the spectacular returns of 1995-99 by increasing estimates of 21st Century returns.

**d) *Unpalatable implications.*** Fama and French (2000) point out that average past U.S. stock returns are so high that they exceed estimates of the return to equity (ROE) calculated for U.S. corporations from accounting data. Thus if one uses average past stock returns to estimate the cost of capital, the implication is that U.S. corporate investments have destroyed value; corporations should instead have been paying all their earnings out to stockholders. This conclusion is so hard to believe that it further undermines confidence in the average-return methodology.

One variation of the average-past-returns approach is worth discussing. One might take the view that average past equity returns in other countries provide relevant evidence about U.S. equity returns. Standard international data from Morgan Stanley Capital International,

---

[2] One crude way to handle this problem is to measure the annualized variance of returns over a period such as 20 years that is long enough for returns to be approximately serially uncorrelated, and then to adjust the geometric average up by one-half the annualized 20-year variance as would be appropriate if returns are lognormally distributed. Campbell and Viceira (2001, Figure 4.2) report an annualized 20-year standard deviation of about 14% in long-term annual U.S. data, which would imply an adjustment of $0.14^2/2=0.010$ or 1.0%.

available since the early 1970's, show that equity returns in most other industrialized countries have been about as high as those in the U.S. The exceptions are the heavily commodity-dependent markets of Australia and Canada, and the very small Italian market (Campbell 1999). Jorion and Goetzmann (1999) argue that other countries' returns were lower than U.S. returns in the early 20th Century, but this conclusion appears to be sensitive to their omission of the dividend component of return (Dimson, Marsh, and Staunton 2000). Thus the use of international data does not change the basic message that the equity market has delivered high average returns in the past.

## 2. Valuation ratios

An alternative approach is to use valuation ratios—ratios of stock prices to accounting measures of value such as dividends or earnings—to forecast future returns. In a model with constant valuation ratios and growth rates, the famous Gordon growth model says that the dividend-price ratio

$$\frac{D}{P} = R - G, \tag{1}$$

where $R$ is the discount rate or expected equity return, and $G$ is the growth rate of dividends (equal to the growth rate of prices when the valuation ratio is constant). This formula can be applied either to price per share and conventional dividends per share, or to the total value of the firm and total cash paid out by the firm (including share repurchases). A less well-known but just as useful formula says that in steady state, where earnings growth comes from reinvestment of retained earnings which earn an accounting ROE equal to the discount rate $R$,

$$\frac{E}{P} = R. \tag{2}$$

Over long periods of time summarized by Siegel (1998), these formulas give results consistent with average realized returns. Over the period 1802-1997, for example, the average dividend-price ratio was 5.4% while the geometric average growth rate of prices was 1.6%. These numbers add to the geometric average return of 7.0%. Over the period 1871-1997 the average dividend-price ratio was 4.9% while the geometric average growth rate of prices was 2.1%, again adding to 7.0%. Similarly, Campbell and Shiller (2001) report that the average P/E ratio for S&P 500 shares over the period 1872-2000 was 14.5. The reciprocal of this is 6.9%, consistent with average realized returns.

When valuation ratios and growth rates change over time, these formulas are no longer exactly correct. Campbell and Shiller (1988) and Vuolteenaho (2000) derive dynamic versions of the formulas that can be used in this context. Campbell and Shiller show, for example, that the log dividend-price ratio is a discounted sum of expected future discount rates, less a discounted sum of expected future dividend growth rates. In this note I will work with the simpler deterministic formulas.

## II. Current Market Conditions

Current valuation ratios are wildly different from historical averages, reflecting the unprecedented bull market of the last 20 years, and particularly the late 1990's. The attached figure, taken from Campbell and Shiller (2001), illustrates this point. (See p. 9) The bottom left panel shows the dividend-price ratio $D/P$ in January of each year from 1872-2000. The long-term historical average is 4.7%, but $D/P$ has fallen dramatically since 1982 to about 1.2% in January 2000 (and 1.4% today).

The dividend-price ratio may have fallen in part because of shifts in corporate financial policy. An increased tendency for firms to repurchase shares rather than pay dividends increases the growth rate of dividends per share, by shrinking the number of shares. Thus it increases $G$ in the Gordon growth formula and reduces conventionally measured $D/P$. One way to correct for this is to add repurchases to conventional dividends. Recent estimates of this effect by Liang and Sharpe (1999) suggest that it may be an upward adjustment of 75 to 100 basis points, and more in some years. Of course, this is not nearly sufficient to explain the recent decline in $D/P$.

Alternatively, one can look at the price-earnings ratio. The top left panel of the figure shows $P/E$ over the same period. This has been high in recent years, but there are a number of earlier peaks that are comparable. Close inspection of these peaks shows that they often occur in years such as 1992, 1934, and 1922 when recessions caused temporary drops in (previous-year) earnings. To smooth out this effect, Campbell and Shiller (2001), following Graham and Dodd (1934), advocate averaging earnings over 10 years. The price-averaged earnings ratio is illustrated in the top right panel of the figure. This peaked at 45 in January 2000; the previous peak was 28 in 1929. The decline in the S&P 500 since January 2000 has only brought the ratio down to the mid-30's, still higher than any level seen before the late 1990's.

The final panel in the figure, on the bottom right, shows the ratio of current to 10-year average earnings. This ratio has been high in recent years, reflecting robust earnings growth during the 1990's, but it is not unprecedentedly high. The really unusual feature of the recent stock market is the level of prices, not the growth of earnings.

## III. Implications for Future Returns

The implications of current valuations for future returns depend on whether the market has reached a new steady state, in which current valuations will persist, or whether these valuations are the result of some transitory phenomenon.

If current valuations represent a new steady state, then they imply a substantial decline in the equity returns that can be expected in the future. Using Campbell and Shiller's (2001) data, the unadjusted dividend-price ratio has declined by 3.3 percentage points from the historical average. Even adjusting for share repurchases, the decline is at least 2.3 percentage points. Assuming constant long-term growth of the economy, this would imply that the geometric average return on equity is no longer 7%, but 3.7% or at most 4.7%. Looking at the price-averaged earnings ratio, adjusting for the typical ratio of current to averaged earnings, gives an even lower estimate. Current earnings are normally 1.12 times averaged earnings; 1.12/35=0.032, implying a 3.2% return forecast. These forecasts allow for only a very modest equity premium relative to the yield on long-term inflation-indexed bonds, currently about 3.5%, or the 3% safe real return assumed recently by the Trustees.

If current valuations are transitory, then it matters critically what happens to restore traditional valuation ratios. One possibility is that earnings and dividends are below their long-run trend levels; rapid earnings and dividend growth will restore traditional valuations without any declines in equity returns below historical levels. While this is always a possibility, Campbell and Shiller (2001) show that it would be historically unprecedented. The U.S. stock market has an extremely poor record of predicting future earnings and dividend growth. Historically stock prices have increased relative to earnings during decades of rapid earnings growth, such as the 1920's, 1960's, or 1990's, as if the stock market anticipates that rapid earnings growth will continue in the next decade. However there is no systematic tendency for a profitable decade to be followed by a second profitable decade; the 1920's, for example, were followed by the 1930's and the 1960's by the 1970's. Thus stock market optimism often fails to be justified by subsequent earning growth.[3]

A second possibility is that stock prices will decline or stagnate until traditional valuations are restored. This has occurred at various times in the past after periods of unusually high stock prices, notably the 1900's and 1910's, the 1930's, and the 1970's. This would imply extremely low and perhaps even negative returns during the adjustment period, and then higher returns afterwards.

The unprecedented nature of recent stock market behavior makes it impossible to base forecasts on historical patterns alone. One must also form a view about what happened to drive stock prices up during the 1980's and particularly the 1990's. One view is that there has been a structural decline in the equity premium, driven either by the correction of mistaken perceptions of risk (aided perhaps by the work of economists on the equity premium puzzle), or by the reduction of barriers to participation and diversification by small investors.[4] Economists such as McGrattan and Prescott (2001) and Jagannathan, McGrattan, and Scherbina (2001) argue that the structural equity premium is now close to zero, consistent with theoretical models in which investors effectively share risks and have modest risk aversion, and consistent with the view that the U.S. market has reached a new steady state.

---

[3] Vuolteenaho (2000) notes, however, that U.S. corporations were unusually profitable in the late 1990's and that profitability has some predictive power for future earnings growth.

[4] Heaton and Lucas (1999) model barriers of this sort. It is hard to get large effects of increased participation on stock prices unless initial participation levels are extremely low. Furthermore, one must keep in mind that what matters for pricing is the wealth-weighted participation rate, that is, the probability that a randomly selected dollar of wealth is held by an individual who can participate in the market. This is higher than the equal-weighted participation rate, the probability that a randomly selected individual can participate.

An alternative view is that the equity premium has declined only temporarily, either because investors irrationally overreacted to positive fundamental news in the 1990's (Shiller 2000), or because the strong economy made investors more tolerant of risk.[5] On this view the equity premium will return to historical levels, implying extremely poor near-term returns and higher returns in the more distant future after traditional valuations have been restored.

It is too soon to tell which of these views is correct, and I believe it is sensible to put some weight on each of them. That is, I expect valuation ratios to return part way but not fully to traditional levels.[6] A rough guess for the long term, after the adjustment process is complete, might be a geometric average equity return of 5% to 5.5% or an arithmetic average return of 6.5% to 7%.

If equity returns are indeed lower on average in the future, it is likely that short-term and long-term real interest rates will be somewhat higher. That is, the total return to the corporate capital stock is determined primarily by the production side of the economy and by national saving and international capital flows; the division of total return between riskier and safer assets is determined primarily by investor attitudes towards risk. Reduced risk aversion then reduces the equity premium both by driving down the equity return and by driving up the riskless interest rate. The yield on long-term inflation-indexed Treasury securities (TIPS) is about 3.5%, while short-term real interest rates have recently averaged about 3%. Thus 3% to 3.5% would be a reasonable guess for safe real interest rates in the future, implying a long-run average equity premium of 1.5% to 2.5% in geometric terms or about 3% to 4% in arithmetic terms.

Finally, I note that it is tricky to use these numbers appropriately in policy evaluation. Average equity returns should never be used in base-case calculations without showing alternative calculations to reflect the possibilities that realized returns will be higher or lower than average. These calculations should include an alternative in which equities underperform Treasury bills. Even if the probability of underperformance is small over a long holding period, it cannot be zero or the stock market would be offering an arbitrage opportunity or "free lunch". Equally important, the bad states of the world in which underperformance occurs are heavily weighted by risk-averse investors. Thus policy evaluation should use a broad range of returns to reflect the uncertainty about long-run stock market performance.

---

[5] Campbell and Cochrane (1999) present a model in which investors judge their well-being by their consumption relative to a recent average of past aggregate consumption. In this model investors are more risk-tolerant when consumption grows rapidly and they have a "cushion of comfort" relative to their minimum expectations. The Campbell-Cochrane model fits past cyclical variations in the stock market, which will likely continue in the future, but it is hard to explain the extreme recent movements using this model.

[6] This compromise view also implies that negative serial correlation, or mean-reversion, is likely to remain a characteristic of stock returns in the 21st Century.

Figure 4. S&P Composite Stock Data, January Values 1872-1997

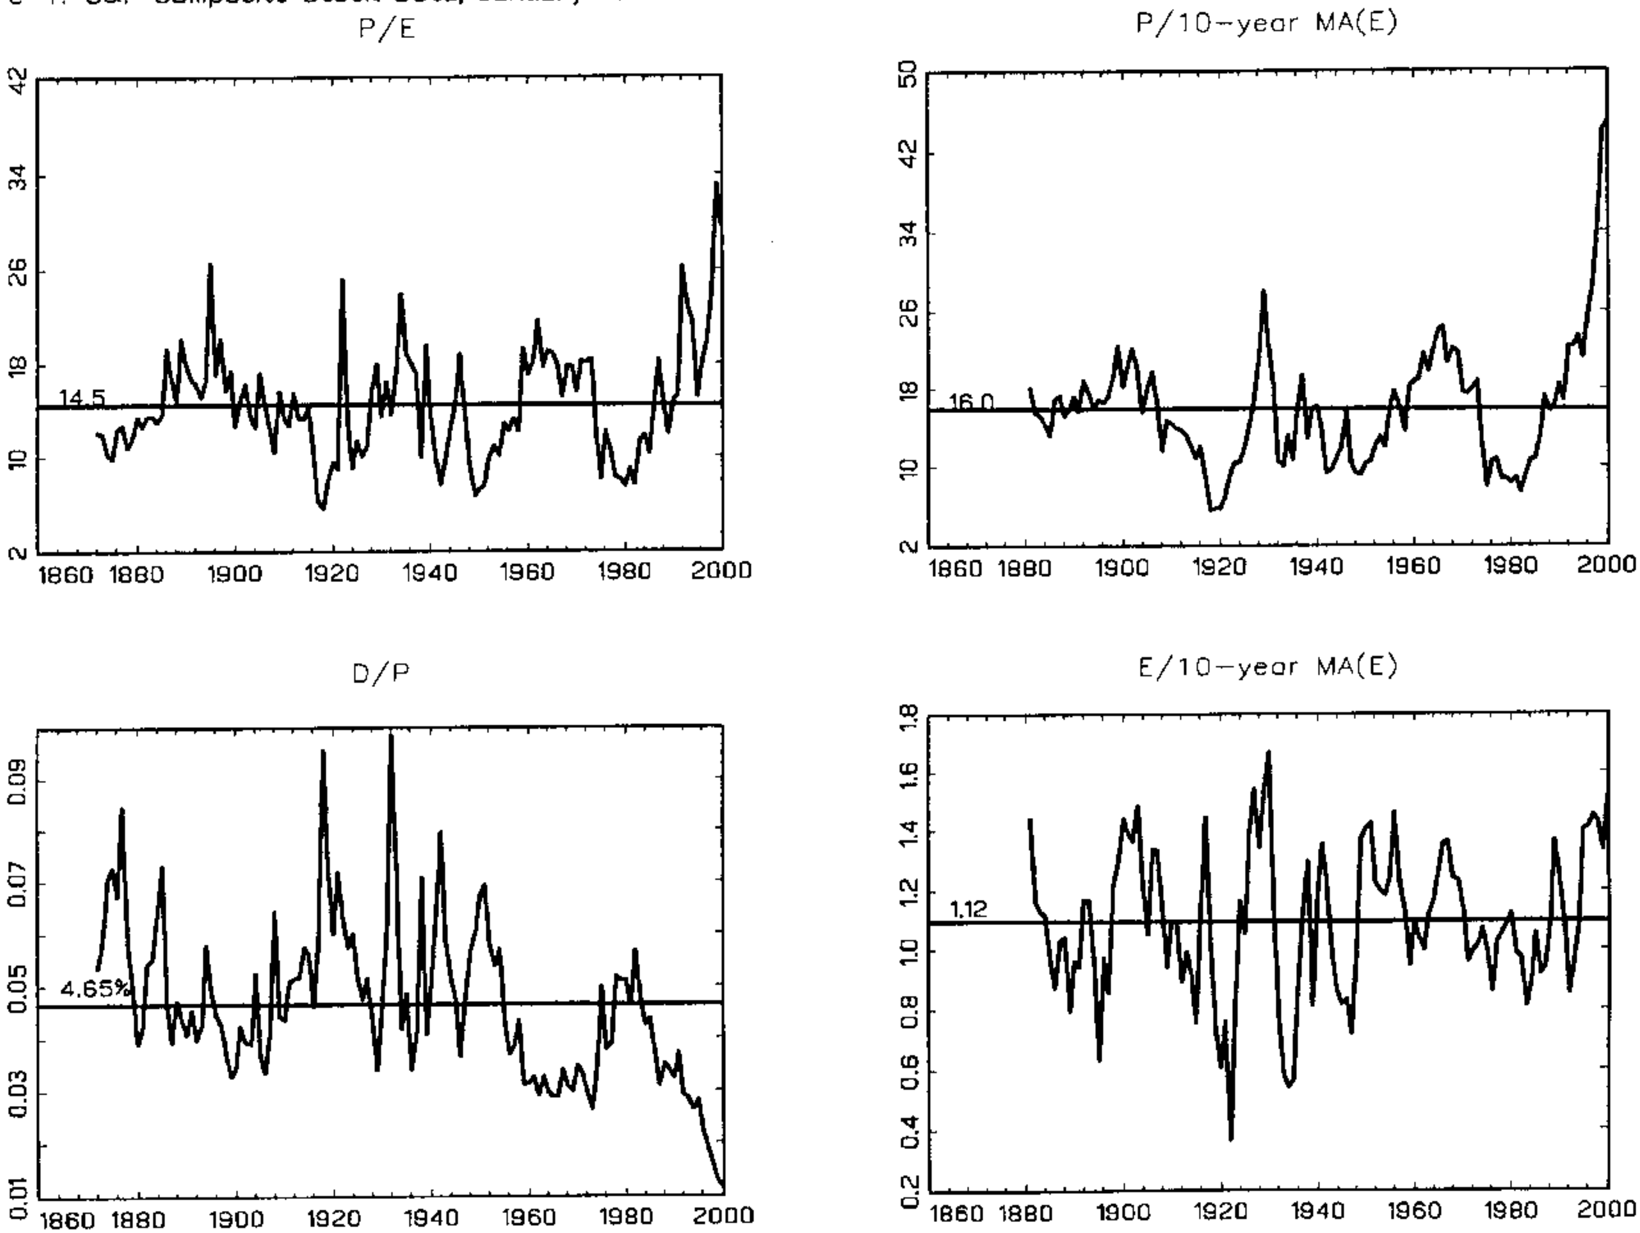

## Bibliography


Campbell, John Y. and John H. Cochrane, 1999, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior", *Journal of Political Economy* 107, 205-251.

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

Campbell, John Y. and Robert J. Shiller, 2001, "Valuation Ratios and the Long-Run Stock Market Outlook: An Update", NBER Working Paper No. 8221.

Campbell, John Y. and Luis M. Viceira, 2001, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, forthcoming Oxford University Press, New York, NY.

Dimson, Elroy, Paul Marsh, and Mike Staunton, 2000, *The Millennium Book: A Century of Investment Returns,* ABN Amro and London Business School, London.

Fama, Eugene and Kenneth R. French, 2000, "The Equity Premium", unpublished paper, University of Chicago and MIT.

Graham, Benjamin and David Dodd, 1934, *Security Analysis*, McGraw-Hill, New York, NY.

Heaton, John and Deborah Lucas, 1999, "Stock Prices and Fundamentals", *NBER Macroeconomics Annual* 213-242, MIT Press, Cambridge, MA.

Jagannathan, Ravi, Ellen R. McGrattan, and Anna Scherbina, 2001, "The Declining US Equity Premium", NBER Working Paper No. 8172.

Jorion, Philippe and William N. Goetzmann, 1999, "Global Stock Markets in the Twentieth Century", *Journal of Finance* 54, 953-980.

Liang, Nellie and Steven A. Sharpe, 1999, "Share Repurchases and Employee Stock Options and their Implications for S&P 500 Share Retirements and Expected Returns", unpublished paper, Board of Governors of the Federal Reserve System.

McGrattan, Ellen R., and Edward C. Prescott, 2001, "Is the Stock Market Overvalued?", NBER Working Paper No. 8077.

Shiller, Robert J., 2000, *Irrational Exuberance*, Princeton University Press, Princeton, NJ.

Siegel, Jeremy, 1998, *Stocks for the Long Run*, 2nd ed., McGraw-Hill, New York, NY.

Vuolteenaho, Tuomo, 2000, "Understanding the Aggregate Book-to-Market Ratio and its Implications to Current Equity-Premium Expectations", unpublished paper, Harvard University.

# What Stock Market Returns to Expect for the Future: An Update

**Peter A. Diamond, Professor of Economics**
**Massachusetts Institute of Technology**
**July 23, 2001**

This note updates the calculations in my previous analysis of this issue (Social Security Bulletin, 2000, vol. 63, no. 2, pp. 38-52).* The calculations address two issues. First, what are the implications of assuming an annual 7% real return on equities throughout the next 75 years (along with the assumptions in the Trustees' Report), as has been the practice in OCACT projections of Social Security reform proposals that include equities. While the numbers are changed some from those based on the end of 1998, calculations done for the end of 2000 and the end of the first quarter of 2001 continue to show that a 7% return throughout the next 75 years from these starting points is implausible.

Second, what are the implications for stock market values in ten years if there is to be a lower rate of return for the next decade, followed by a return to the historical average return thereafter. As before, the returns over the next decade need to be very low, indeed an unchanged nominal value for stocks at the end of the decade is roughly consistent with close to a 7% return thereafter.

The calculations reported here are based on the Gordon formula, relating stock values to returns and the growth of returns. A first step in considering stock market returns is to project the future net cash flow to stockholders. This is normally done in three steps. First is to estimate the current net cash flow. Second is to adjust that for reasons to believe that the long-run relationship to GDP may be different from the current relationship. And third is to assume a constant relationship to GDP given the first two steps.

The cash flow to holders of publicly traded stocks as a whole contains many pieces. Easy to measure is the flow of dividends. Then there is the cash flow arising from share repurchase. This happens in two ways – direct repurchase of a corporation's own shares and acquisition of the shares of other corporations for cash or debt. Sometimes acquired shares are retired and sometimes they are not. This may be a complication in estimation given how data are presented – I have not reviewed measurement in data sources.

In order to maintain any given fraction of the value of shares outstanding, there are also pieces that are equivalent to negative cash flows. When employees exercise stock options and so acquire shares at less than market value, there is a dilution of the stock value of existing owners. This can be approached by thinking about the excess of market value over exercise price or by considering the value of options that are given to employees.

---

* See article beginning on p. 17.

I am grateful to Mauricio Soto for excellent research assistance, doing the calculations reported here. I am also grateful for financial support from the Retirement Research Center at Boston College.