Economic Review

January/February 1995
Volume 80, Number 1

Federal Reserve Bank of Atlanta

In This Issue:

Economic Review

January/February 1995, Volume 80, Number 1

Economic Review

Federal Reserve Bank of Atlanta

The *Economic Review* of the Federal Reserve Bank of Atlanta presents analysis of economic and financial topics relevant to Federal Reserve policy. In a format accessible to the nonspecialist, the publication reflects the work of the Research Department. It is edited, designed, produced, and distributed through the Public Affairs Department.

Views expressed in the *Economic Review* are not necessarily those of this Bank or of the Federal Reserve System.

Free subscriptions and limited additional copies are available from the Public Affairs Department, Federal Reserve Bank of Atlanta, 104 Marietta Street, N.W., Atlanta, Georgia 30303-2713 (404/521-8020). Change-of-address notices and subscription cancellations should be sent directly to the Public Affairs Department. Please include the current mailing label as well as any new information. ISSN 0732-1813

Contents

Federal Reserve Bank of Atlanta *Economic Review*
January/February 1995, Volume 80, Number 1

Given the vast sums of money that mutual and pension fund managers invest, an important question is how they should go about deciding which assets, especially which stocks, they should purchase. One school of thought argues that investment policies should reflect some set of social values. This study examines three questions about the financial implications of effective socially responsible investing in common stocks—that is, socially responsible investment intended to change firms' behavior.

The first question concerns what socially responsible investors can do to effectively influence firms' investment policies. The second question is, under what conditions, if any, will the securities markets permit effective socially responsible investment? Third, what impact will socially responsible investment have on the performance of portfolios that follow it?

The analysis has two implications for fund managers and investors who want to change firms' behavior. The first is that the investment strategy should focus on buying shares of small socially responsive firms. The second is that investors who owned targeted socially responsible stocks before socially responsible investment began will realize above-market rates of return in the short run; however, once socially responsible investors stop bidding up the price, investors will receive reduced returns.

Although current inflation rates are relatively benign, the costs of unexpected inflation, even at low rates, remain substantial for individual firms and consumers. Many types of planning decisions, such as businesses' and governments' plans for expected expenses and revenues, hinge on inflation forecasts.

This article provides an overview of the effects of inflation and the significance of inflation forecasting. The author first considers how forecasting models are specifically designed to fill the needs of particular users. The analysis examines two statistical models—the Phillips curve and money demand/monetarist models—that employ standard economic theory to suggest variables that help predict inflation. Forecasts from a simple, but widely used, version of each model are then compared with simple time series models that include only past data on inflation. This comparison using standard accuracy criteria shows that the economic models did not perform much better than the simplest time series forecasting model. The author concludes that future research should focus on estimating dynamic models that are by design more structural and that may help uncover the sources of inflation.

Douglass North earned a share of the 1993 Nobel prize for economics for two decades of research that culminated in the development of an innovative framework for analyzing economic history. This review essay discusses the book that most comprehensively presents North's paradigm, which characterizes history as the record of an evolving game in which institutions, organizations, and individuals function as the rules, teams, and players. Through examples, the reviewer illustrates how North's game theoretic paradigm can serve not only as a tool for analyzing historical events but also as a methodological bridge between the diverse branches of the social sciences and humanities.

Some Lessons from Basic Finance for Effective Socially Responsible Investing

Larry D. Wall

Mutual funds and pension plans are two of the most rapidly growing types of financial intermediaries in the United States. Together they accounted for almost 40 percent of all financial intermediaries' assets at the end of 1993 (see George G. Kaufman and Larry R. Mote 1994). Given the vast sums of money that mutual and pension fund managers invest, an important question is how they should go about deciding which assets, especially which stocks, they should purchase. Funds that invest in stocks have generally been actively managed—that is, the fund manager buys stocks he or she believes are undervalued and sells stocks thought to be overvalued. This approach to investing has been challenged in recent years, however.

Modern finance theory presents one challenge, arguing that funds should buy widely diversified portfolios whose composition roughly approximates the entire stock market. This position is based on the efficient markets argument, which holds that stocks are correctly priced and investors cannot systematically find stocks that are either under- or overvalued. From this perspective, active management of stocks only wastes value by increasing transactions costs.[1]

The other challenge to traditional active management comes from those who argue that investment policies should reflect some set of social values.[2] Advocates of this point of view argue that socially responsible investing may help improve the world. Some would also argue that socially responsive corporations may be more profitable in the long run because their

The author is a research officer in the financial section of the Atlanta Fed's research department. *He thanks Peter Abken, Andy Krikelas, and Pam Peterson for helpful comments.*

workers are likely to be more productive and because changing government regulations and social pressures are less likely to affect such corporations adversely.

Investment based on social goals falls under the broad category of socially responsible investment.[3] Given that different people place different priorities on various social goals, there is no single, universally accepted definition of social responsibility. In practice, socially responsible investment criteria work by limiting the universe of stocks that a fund manager considers via some combination of negative (exclusionary) and positive (inclusionary) screens. For example, participation in the tobacco and weapons industries might constitute "socially undesirable" activities included in negative screens; "socially desirable" activities allowed under positive screens might include following good environmental practices and seeking to have good employee relations.

While the economic merits of active versus passive management of funds have been a source of continuing debate among academics and practitioners, the economic implications of socially responsible investing have received somewhat less attention. It is likely that concerning this issue the gap between academic opinions and practitioner views is somewhat larger. This study seeks to close that gap by focusing on what will here be called effective socially responsible investment in common stocks—that is, the buying (selling) of publicly traded common stocks on the basis of some social criteria in order to increase (decrease) the firms' level of investment in plant and equipment. In other words, the discussion analyzes the implications of socially responsible investment policies that a mutual fund, pension fund, or individual investor might follow. In particular, this study addresses three questions: (1) What must socially responsible investors do to effectively influence firms' investment policies? (2) Under what conditions, if any, will securities markets permit effective socially responsible investment? (3) If effective socially responsible investment is possible for at least some securities, how will including such securities affect the performance of portfolios? The analysis is largely based on fundamental finance principals covered in standard undergraduate finance texts; in particular this study will use the terminology and notation of Pamela P. Peterson's (1994) introductory text.[4]

The results of this study will be most relevant to investors with social concerns who want to change firms' behavior. This discussion will indicate both when significant social changes are likely to result from investment decisions and what the implications of effective socially responsible investment will be for an investor's rate of return. The analysis will be far less valuable for investors whose primary concern is matching their investments to their personal values rather than serving larger social goals. Existing finance theory, which generally assumes that investors' only concern is the risk and return from their asset holdings, has little to say about optimal investment policies for investors who also want to take into account the way in which firms obtain their earnings. The analysis will also have limited value for investors who plan on buying securities that are not publicly traded and who therefore face issues very different from those discussed below.[5]

Existing Studies of Socially Responsible Investment

Efficient Market Analysis. Financial economists usually begin any analysis of stock markets with the efficient markets hypothesis, which holds that stock prices reflect all publicly available information.[6] In this context of what is called "semistrong efficiency," no investor may use publicly available information to obtain an above-market rate of return without a commensurate increase in risk.[7] Socially responsible investment portfolios must earn a market rate of return after adjusting for their level of (nondiversifiable) risk.

Arguments for Socially Responsible Investment. Advocates of socially responsible investment attack the efficient markets perspective in two ways. First, they argue that a large and growing body of evidence suggests that markets are not semistrong efficient. Maria O'Brien Hylton (1992), for example, points to a number of studies that suggest inefficiency in financial markets. According to Hylton these studies imply that socially responsible investors could profit from market inefficiency resulting from speculation.

Second, some advocates of socially responsible investment argue that, given the inefficiency of the market, "socially responsible" firms may generate above-market returns and "socially irresponsible" firms may generate below-market returns (for example, see Severyn T. Bruyn 1987). For instance, firms that follow socially irresponsible environmental policies may face increased costs and reduced market share as governments and consumers hold corporations to ever-increasing standards of environmental cleanliness. In contrast, firms that follow environmentally friendly policies may have increasing market opportunities around the world as people become more concerned about their environment.

Counterarguments. Although the claim that markets have been found to be inefficient has substantial support, it is not strong enough to provide a compelling argument in support of socially responsible investment. Numerous studies have found evidence that markets are in fact efficient in a wide variety of circumstances.[8] It is therefore a mistake to generalize from the findings of the few studies that suggest market inefficiency to conclude that market prices are consistently inaccurate indicators of a firm's value.

Moreover, the case for socially responsible investment would not be proved even if it were shown to be true both that (1) financial markets are inefficient and (2) socially responsible firms are more profitable than irresponsible firms.[9] The source of the market inefficiency could be unrelated to whether firms are socially responsible, in which case an investor following socially responsible investment policies would not earn a higher rate of return. Another possibility is that firms that generate high profits are more likely to be overvalued by the stock market in which investors following socially responsible investing may earn below market rates of return.

Empirical Evidence. The theoretical arguments that socially responsible investment can generate superior returns may be sufficient to leave some doubt about the validity of the efficient markets view but are not strong enough to make a persuasive case that socially responsible investors will in fact earn superior returns. The obvious next step is to look at the empirical evidence concerning how existing socially responsible investment funds perform relative to the market. A simplistic approach to analyzing the evidence yields mixed results. Some socially responsible investment mutual funds over some periods of time produce rates of return in excess of other mutual funds and the market.[10] However, other comparisons of socially responsible investment funds find that they produce inferior returns over some periods of time.[11]

Thus, finding an answer to the question of whether socially responsible investing outperforms conventional investing is not an easy task. Two problems in particular raise doubts about simplistic comparisons of socially responsible investment mutual fund results. First, stock returns contain an element that is not predictable: that is, random luck will influence the results. For example, unexpected changes in consumer tastes may turn out to favor industries with historically good (or bad) employee relations. Thus, a socially responsible investment fund that focuses on firms with good employee relations may earn higher (or lower) returns for reasons totally unrelated to the fund's investment criteria.

A second problem in making comparisons is that riskier stocks should perform differently than less risky stocks. The performance of riskier stocks is likely to be better during some periods and worse during others relative to less risky stocks. Further, riskier stocks should provide higher returns over the long run to compensate investors for the greater risk.[12] Thus, any analysis of the performance of socially responsible investment funds should take account of risk.

Sally Hamilton, Hoje Jo, and Meir Statman (1993) analyzed the performance of socially responsible investment mutual funds relative to other mutual funds using a methodology designed to take account of random luck and differences in the riskiness of various portfolios. The measure they use, Jensen's alpha, is based on the capital asset pricing model (CAPM).[13] They found that socially responsible investment mutual funds started prior to 1986 outperform comparable conventional mutual funds and that socially responsible investment mutual funds started in 1986 or later underperform conventional mutual funds. However, in neither case are the results statistically significant, which means that the differences in performance could be due to random luck. Thus, their results suggest that socially responsible investment mutual funds' performance is neither better nor worse than funds that do not follow socially responsible investment principles. The Hamilton, Jo, and Statman results are consistent with the prediction of the efficient markets hypothesis, which suggests that no portfolio selection criteria, whether based on social values or a particular economic theory, will produce consistently superior investment results.

For socially responsible investment to make a real difference in the world changes must begin with firms' investment policies.

Potential socially responsible investors may be disappointed that existing evidence does not show that such a strategy produces superior returns. They may take comfort, however, from the evidence that neither does it produce inferior returns. The next question facing potential socially responsible investors is whether they can help improve the world through their investments while obtaining portfolio returns comparable to conventional funds on a risk-adjusted basis.

Effective Socially Responsible Investment

For socially responsible investment to make a real difference in the world changes must begin with firms' investment policies. Socially desirable firms must be encouraged to expand by investing more in plant and equipment than they would without socially responsible investment, and not investing in socially undesirable firms, at least as long as they follow undesirable policies, should force them to reduce their level of investment. In order to understand how purchases and sales of stock motivated by socially responsible investment can influence investment policy two issues need to be considered: (1) How do firms decide whether to invest in a project? And (2) how could socially responsible investment change the results of the investment analysis?

Firms that maximize shareholder wealth should seek to invest in projects whose returns equal or exceed the rate of return required by investors.[14] The preferred method used to determine whether the project returns are sufficiently high is called net present value (NPV).[15] The formula for calculating net present value and an example are presented in Box 1. Essentially, net present value works by comparing the initial investment in a project with the adjusted value of future cash flows. The adjustment procedure uses a discount factor to reduce the value of future inflows to compensate for the time value of money and the risk of the project. A large reduction in the value of future cash inflows—a high discount factor—discourages firms from investing in new projects.[16] Conversely, a low discount factor will encourage investment.

This discussion suggests that socially responsible investment might change firms' investment practices if it could change the discount factor applied to investment decisions. Can socially responsible investment change the discount factor? The answer to this question turns out to be yes: if socially responsible investment can increase a firm's stock price, it may decrease the discount factor. Changes in stock price influence the discount factor because the discount factor should depend on a firm's cost of funds. Changes in a firm's stock price influence the cost of equity to the firm and, hence, its estimated cost of funds for future periods (as measured by the weighted average cost of capital [WACC]). A higher stock price implies a lower discount factor applied to future cash flows and hence greater investment by the firm. Following the socially responsible investment prescription of investing in socially desirable firms could therefore induce these firms to increase their investment in new and expanded projects. On the other hand, reducing a firm's stock price may drive up the discount factor. The firm's higher WACC would discourage further investment.

One caveat to this analysis should be noted. In order for socially responsible investment to cause a change in a firm's investment policy the change in the stock price must be large and must persist over an extended period of time. NPV calculations are subject to substantial measurement error in at least three areas: (1) the estimation of the expected cash flows, (2) the estimation of the cost of equity, and (3) the process of adjusting for differences between the risk of the project and the firm's average risk. Firms generally recognize these measurement errors and try to offset them by using rules that are somewhat more conservative than is implied by a straightforward calculation of NPV using WACC. For example, firms may not calculate a new WACC daily to reflect changes in their stock price. Thus, small, temporary changes in share prices are unlikely to cause any change in a firm's investment policy.

Socially Responsible Investment and Common Stock Prices

The above analysis suggests that socially responsible investment may be effective in changing firms' investment policies if it can cause a substantial, long-run change in their stock prices. Thus, the question becomes that of whether socially responsible investment can change a firm's stock price.

An efficient markets perspective says that a firm's stock price reflects all available information about a stock and that all stocks earn only a rate of return commensurate with their market risk. Thus, under efficient markets, socially responsible investment must be ineffective. Yet, as anyone learns in one semester of economics, market prices are set by supply and demand. Since socially responsible investment policies change either the supply or demand for a stock, the law of supply and demand suggests that the price should also change. Resolving this seeming conflict

Box 1
Project Valuation

Net Present Value

Firms that maximize shareholder wealth should seek to invest in projects whose returns, adjusted for the time value of money and risk, equal or exceed the rate of return required by investors. The method used to adjust projects' earnings for the time value of money and risk is called net present value (*NPV*).

$$NPV = \sum_{t=0}^{T} \frac{CF_t}{(1+r)^t},$$

where NPV = net present value, T = useful life of the project, t = number of discounting periods, CF_t = cash flow at the end of period t, and r = cost of capital (see Peterson 1994, 399-405).

The method of calculating *NPV* may be illustrated by applying it to the three hypothetical projects given in Table 1.[1] For convenience, assume that all three projects are of equal risk and that their risk is also equal to the average risk of the firm Peachtree Multimedia Software.[2] In all three cases the expected cash flows over the life of the projects exceeds the investment cost. However, before investing in any of the projects Peachtree Multimedia Software should determine whether the expected cash flows adequately compensate for the time value of money and the riskiness of the projects. The net present value calculation takes account of the time value and riskiness by discounting future cash flows at the rate r. Suppose, for example, that Peachtree Multimedia Software decides the appropriate discount factor is 10 percent.[3] In this case the *NPV* of project A is calculated as follows (with rounding to the nearest dollar):

$$NPV = -\$10{,}000 + \$4{,}000/1.10 + \$4{,}500/(1.10)^2 + \$4{,}100/(1.10)^3 = \$187.$$

Project A would be accepted in this case because it has a net present value of $187, meaning that the value of its expected cash inflows adjusted for risk and the time value of money exceeds the required investment by $187. In finance terms, project A should be undertaken because it is a positive *NPV* project. The net present values of projects B and C are given in the second and third columns of Table 2. The adjusted value of the expected cash inflows from these is less than the value of the required outlay. These projects have a negative *NPV* and should be rejected.

Suppose that Peachtree Multimedia Software is considered to be socially responsible and socially responsible investors would like to see the firm grow more rapidly. Stock investors may not be able to change the cash flows to the project, but they may be able to change the discount factor applied to the cash flows. For example, suppose that the discount factor is lowered to 8 percent. Table 2 shows that in this case project A retains a positive *NPV*, project B gains a positive *NPV*, but C remains negative. Thus, lowering the discount factor has resulted in additional investment and expanded the size of the firm. Conversely, suppose that Peachtree Multimedia Software was considered to be socially irresponsible and socially responsible investors would prefer that it make no new investment. In this case an increase in the discount factor to 12 percent would make all of the projects have a negative *NPV*.

Weighted Average Cost of Capital

The discussion of net present value suggests that determining the appropriate discount factor is important in project valuation. The appropriate discount factor depends on the firm's cost of funds. Calculating the cost of funds would be a trivial process for a firm that was all debt-financed if there were no taxes. In this case the cost of funds would be the interest rate on the debt. In a firm subject to taxes with a mixture of debt and equity the measurement becomes more complicated. A firm that seeks to measure its cost of capital must address three issues: (1) the cost of its equity, (2) how it can adjust for the differences in the tax treatment of debt and equity, and (3) how it can adjust for the proportion of debt and equity contained in its capital structure.[4]

There are a variety of ways to calculate a firm's cost of common equity.[5] One widely used model is a dividend growth valuation model.[6] The dividend growth model says that the value of a stock should be the net present value of its expected dividend payments. Further, the version of the model used makes the simplifying assumption that dividends are expected to grow at a constant rate forever.[7] For example, suppose that Peachtree Multimedia Software is currently paying dividends of $7, these dividends are expected to grow at a 2 percent annual rate, and the firm's stock price is $50:

$$r_e = D_t/P_0 + g = 7/50 + 2\% = 16\%,$$

where r_e = required return on equity, D_t = dividends per share of stock at time t, P_0 = present value of a share of stock = its price at time zero, and g = expected growth rate of dividends per share. Thus, the required rate of return on Peachtree Multimedia Software stock is 16 percent.

Interest payments on debt may be deducted by firms whereas dividend payments to stockholders are not tax deductible.[8] Thus, to determine the cost of the two sources of funds to the firm, an adjustment must be made

continued on next page

continued from page 5

to recognize the difference in tax treatment. The standard approach is to reduce the interest rate the firm pays to reflect the fact that the government shares in the cost of making the interest payment (in the same sense that the government shares in the firm's revenue and its other costs of production such as wage payments to workers). Suppose, for example, that Peachtree Multimedia Software could issue new debt that yielded 8 percent per year to investors and that Peachtree was in a 25 percent tax bracket. In this case the after-tax cost of debt to Peachtree is calculated as follows (see Peterson 1994, 636-37):

$$r_d^* = r_d(1 - \tau) = 8\%(1 - .25) = 6\%.$$

Thus, the after-tax cost of debt to Peachtree Multimedia Software is 6 percent.

In general the average after-tax costs of debt and equity will differ. Thus, if a firm's capital structure contains a mix of debt and equity, these costs need to be averaged to determine the firm's overall cost of capital. The weighted average cost of capital (*WACC*) calculates this figure, as its name suggests, by taking weighted averages of the cost of debt and equity where the weights equal the proportion of funding that the firm obtains from each source. If Peachtree Multimedia Software is funded by 60 percent debt, then its *WACC* may be calculated as

$$WACC = w_d r_d + w_e r_e = (60\% \bullet 6\%) + (40\% \bullet 16\%) = 10\%,$$

where w_d = proportion of the firm's funding obtained from debt and w_e = proportion of the firm's funding obtained from equity.[9] Thus, Peachtree's *WACC* is equal to 10 percent.[10]

Table 1
Project Cash Flows

	Project A	Project B	Project C
Initial investment outflow	$10,000	$10,000	$10,000
Net cash inflows			
Year 1	$4,000	$3,700	$3,600
Year 2	$4,400	$4,200	$4,000
Year 3	$4,000	$4,050	$4,000
Net cash flows	$2,400	$1,950	$1,600

Table 2
Net Present Values

	Project A	Project B	Project C
NPV given			
initial discount factor = 10%	$187	($122)	($416)
reduced factor for socially responsible firm = 8%	$559	$242	($62)
increased factor for socially irresponsible firm = 12%	($163)	($466)	($750)

Notes

1. The formula for calculating net present value is widely accepted, but there are important questions about its implementation. For example, how should cash flows from the project be measured? And how should the value of the implicit option to defer the project be valued? Given the focus of this article, the issue of how to measure cash flows may be especially important in some cases. For example, failure to incorporate future outlays by the firm to clean up environmental contamination will bias project selection in favor of environmentally unsound projects. This research assumes that these measurement issues have been correctly addressed by individual firms, however, since the issue of how to measure both cash flows and the option of deferral for individual projects are management issues that are generally outside the control of investors.
2. These assumptions about project risk are not necessary for the conclusions of this section. They are used only to simplify the analysis.
3. A discussion of how to calculate r will follow in the next section. Note that in some cases a time varying discount rate may be more appropriate. For example, r may equal 10 percent for the first period and 11 percent for the second period. However, introducing this complication would not change the conclusions of the analysis.
4. The calculation of the cost of capital may also be influenced by changing the mix of debt and equity. As a part of the assumption that management maximizes shareholder wealth, this article assumes that firms use the lowest cost mix of debt and equity (referred to as optimal capital structure).
5. In order to simplify the analysis Peachtree Multimedia Software is assumed to have no outstanding preferred stock.
6. Another way of calculating the market's expected rate of return on a stock is by estimating the capital asset pricing model (CAPM). In theory these methods should yield similar results. In practice the different methods may not yield identical results, in part because of problems obtaining exact values for some key parameters. This example is based on the dividend growth model because the relationship between stock price and required rate of return is easier to see in a dividend growth model. Also, the CAPM may not be usable for analyzing the impact of socially responsible investment since that model assumes that stock prices are set in efficient markets.
7. This assumption could easily be relaxed without changing the conclusions.
8. This general statement is subject to some exceptions that are not important for the purposes of this study.
9. Note that these weights are typically based on the book value of outstanding debt and equity, but market values could be used instead. If book values are used, then retained earnings are considered part of owner's equity.
10. The formula for *WACC* comes from Peterson (1994, 653). The one difference is that the term for preferred stock is dropped from the equation because Peachtree Multimedia Software is assumed to have no outstanding preferred stock ($w_p = 0$ in Peterson's terminology).

between efficient markets and the law of supply and demand will show whether socially responsible investment may influence the stock price and, if so, when.

The key to reconciling the two views is to note that supply and demand are, in economists' terms, perfectly elastic in efficient markets.[17] That is, any imbalance in supply or demand that threatens to cause the market price to deviate from the firm's underlying value will produce an immediate change in the behavior of other investors. For example, if an excess of demand by socially responsible investors occurs that threatens to push the stock price above its correct value, other investors will reduce their demand for the stock or increase their supply to keep it from exceeding its value and becoming undesirable for non-socially responsible investors. They would rather put their funds in stocks whose market price value is no more than its intrinsic economic value. The net result is that the stock price will remain unchanged. Conversely, an increase in supply from socially responsible investors dumping socially undesirable firms will lead other investors to increase their demand or decrease supply so that the price is unchanged.[18]

Few financial economists would contend that stock prices are perfectly elastic.[19] Hence, the efficient markets view is unlikely to hold exactly for all common stock at all times. However, the market for some stocks is huge, especially those of large firms such as the ones appearing in the often-cited Dow Jones Industrial Average. These stocks are followed by a large number of analysts so that investors have considerable information on these firms. Moreover, trillions of dollars in international investment funds are prepared to move into or out of these stocks if they deviate significantly from the market's perception of their value. Thus, socially responsible investors are likely to have a difficult time influencing the investment policies of very large firms because they are unlikely to cause significant long-run deviations of their stock price from their economic value.

The best chances for socially responsible investors to influence firms' investment through their share price may be by purchasing shares in smaller companies that follow socially desirable policies.[20] Smaller companies are followed by fewer investors, and these investors often have fewer resources at their disposal.[21] Further, socially responsible investors are more likely to be successful in bidding up the price of desirable stocks for extended periods of time than they are in forcing down the price of undesirable stocks over long periods.[22] Investors may try to force a stock price down by selling their shares, but the effect on the stock price may not be significant unless they own a large block of shares.[23] Even if the stock price declines significantly, the drop is likely to be temporary as non-socially responsible investors increase their purchases of what they would perceive as an "undervalued" asset.[24]

Although most fund managers do little short-selling, an extremely aggressive socially responsible investor or fund manager could try to depress an undesirable stock's price via short-selling. Short-selling occurs when investors who do not own a stock borrow shares from an existing owner to sell. An investor who sells a stock short may profit if the stock price subsequently falls and the borrowed shares can be replaced at a lower price, but the short-seller stands to lose if the stock price appreciates. Short-selling to exploit overvalued stocks is generally more costly than buying a stock to exploit undervalued shares for two reasons: (1) short-sellers must pay any dividends earned to the owner of the stock they borrowed, and (2) short-sellers face restrictions on the use of the proceeds of their short-sale to protect the owner of the borrowed shares. Socially responsible investors may find these costs excessive in trying to use short-selling to permanently depress the stock price of undesirable firms.

Implications of Effective Socially Responsible Investment for Stock Returns

A review of existing studies revealed two points of view on the returns to socially responsible investors: (1) the efficient markets view, which states that socially responsible investors should expect a market rate of return adjusted for risk, and (2) a pro-socially responsible investment view, which holds that socially responsible firms will generate returns superior to irresponsible firms over the long run. However, neither of these views explicitly takes account of the impact of effective socially responsible investment on a firm's stock price and investments. If socially responsible investment policies were to be effective in changing a firm's investment policies, there would be clear implications for the returns earned by socially responsible investors.

As Box 2 shows, in the short run, effective socially responsible investment would be good for investors who owned the socially responsible stock before investment boosted prices and who sold as soon as the price increase driven by socially responsible investment stopped. The initial boost in the stock price enables investors to sell their stocks at a higher price and thus earn a higher rate of return. However, in the long run stock returns are lowered. The higher stock price by itself implies an (eventually) lower rate of return when the flow of new socially responsible investment money to the stock slows or stops; a reduction in new purchases of socially responsible investment stocks will slow or stop the rate of appreciation in the firm's stock price. Further, the percentage of return from dividend payments will be reduced as the same dollar value is divided by a larger share price.

Conclusion

This study examines three questions about the financial implications of effective socially responsible investing in common stocks—that is, socially responsible investment intended to change firms' behavior. The first question concerns what socially responsible investors can do to effectively influence firms' investment policies. The results suggest that to increase investment by a firm socially responsible investment must lower the firm's weighted average cost of capital by significantly increasing the firm's stock price for a period of time. Conversely, investors seeking to reduce a socially irresponsible firm's investment must significantly reduce the firm's stock price for a time.

The second question is, under what conditions, if any, will the securities markets permit effective socially responsible investment? The analysis suggests that although financial markets may not be perfectly efficient, socially responsible investors are likely to find it difficult to cause significant, long-run changes in a firm's stock price. They are more likely to influence the stock price of smaller firms and to be more effective in raising stock prices than in lowering them.

Third, what impact will socially responsible investment have on the performance of portfolios that follow it? Effective socially responsible investment will generate above-market rates of return for investors who owned targeted socially responsible stocks prior to the beginning of socially responsible investment. However, investors in socially responsible investment stocks receive reduced returns once a stock becomes a mature socially responsible investment, mature in the sense that no new funds are flowing into the stock.[25]

This analysis has two implications for fund managers and investors who want to change firms' behavior. First, the investment strategy should focus on buying shares of small socially responsive firms. Second, to the extent that this strategy influences share prices it may earn above-market rates of return in the short run, but it is likely to earn below-market rates in the long run. One issue not discussed above is the risk implications of such a focused investment strategy. John H. Langbein and Richard A. Posner (1980) point out that socially responsible investment policies (effective or ineffective) may increase the riskiness of the portfolio by limiting its diversification. Further, a focus on the stocks of small firms may also increase risk.

Some analysts have argued that mutual and pension funds should be subjected to some sort of social requirements along the lines of those the Community Reinvestment Act imposes on banks.[26] This article has two implications for the question of imposing socially responsible investment requirements on funds holdings of stock. The first is that, unless properly designed, such requirements may have little impact on firms' behavior. Further, to the extent that socially responsible investment requirements are effective in influencing firms, these requirements may reduce the level of returns obtained by pension and mutual funds. Thus, before any such requirements are imposed, both the social benefits of the requirements and the desirability of imposing a de facto tax on savings should be considered.

Box 2

Implications of Effective Socially Responsible Investment for Investors

The impact of effective socially responsible investment on investor returns may be seen in an example involving a hypothetical firm, We Care Widgets. We Care Widgets produces expected profits of $6 per share, all of which are paid as stockholder dividends at the end of each period. The firm has no debt, and it is assumed that its earnings are equal to their expected value throughout this example. It is also assumed that the required rate of return by non-socially responsible investors in the company will be constant throughout the period.[1] The stock of We Care Widgets sells for $30 at both time 0 and at time 1 (the end of the first period). However, immediately after the end of the first period, the firm is recognized as being outstanding in meeting its social responsibilities, and socially responsible investors begin purchasing the stock. The stock price increases to $55 at time 2, and it further increases to $60 at time 3. However, after time 3 the flow of additional socially responsible investment funds into the stock ends, and the stock price remains at $60 at time 4.

The measurement of return on a stock should take account of both cash received by the investor (dividends or interest) and any change in the market value of the asset:

$$\textit{Return on a stock} = \frac{P_1 - P_0}{P_0} + \frac{D_1}{P_0},$$

where P_0 = stock price at time 0, P_1 = stock price at time 1, and D_1 = dividends received at the end of period.[2] Thus, the holding period return on We Care's stock for the first period is $HPR_1 = 6/30 + (30 - 30)/30 = 20\%$. Given our assumptions, this 20 percent is also the *WACC* of We Care Widgets.

The increase in demand from socially responsible investment funds that boosts We Care's stock price to $55 during the second period results in a rather substantial holding period return of 103.3 percent for the second period, which is calculated as follows: $HPR_2 = 6/30 + (55 - 30)/30 = 103.3\%$. Thus, any investors who purchased the stock immediately before time 1 and who sold immediately after time 2 would have more than doubled their investment in a single period.

The slowdown in new socially responsible investment after time 2 reduces the rate of return to investors. The holding period return for the third period drops to 20 percent, in this case exactly the same rate of return earned prior to socially responsible investors' discovery of We Care Widgets. Moreover, when the flow of net new funds into the company's stock stops during period 4, the holding period rate of return to investors falls further to 10 percent: $HPR_4 = 6/60 + (60 - 60)/60 = 10\%$. Assuming that socially responsible investors stabilize the price of We Care Widgets at $60 thereafter, future holding period returns will remain at 10 percent.

Notes

1. All of the assumptions in this example are made to simplify the discussion. Similar results could be shown to hold under more general assumptions.
2. See Peterson (1994, 260-65) for a discussion of measuring the return on a stock.

Notes

1. Fierman (1994) argues that the management fees and trading costs incurred by many pension and mutual fund managers cannot be justified by their funds' performance. Her analysis, which looks at money managers in general, finds that only 26 percent of the 2,700 managers of equity funds for pensions exceeded the S&P 500 over the last ten years.
2. See Kinder (1993) for a short overview of social investing.
3. The allocation of funds to social causes via socially responsible lending is only one of a number of ways that investment resources may be allocated in "socially sensitive" ways. See Srinivasan (1994) for a discussion of development lending where an information asymmetry exists between the borrower and the lender. Also see Tschinkel and Wall (1994) for a discussion of investments that generate public gains that can be (partially) captured by the government and shared with the private developer.
4. In a recent study Hamilton, Jo, and Statman (1993, 66) allude to this study's answers to the first and third questions. However, their article does not explain the logic behind their arguments, presumably because underlying theory was assumed to be obvious to the finance practitioners for whom the publication is intended. A recent working paper by Knoll (1994) reaches the same conclusions to questions 1 and 3. His paper is somewhat more technical and provides a more exhaustive discussion of the issues raised by socially responsible investing.
5. For example, buyers of corporate debt yielding a below-market rate need to be assured that their investment will go to expand the firm's operations rather than to expand the shareholders' returns.
6. This form of the efficient markets hypothesis, which is called the semistrong form of efficiency, is probably the most widely applied. The two other forms of market efficiency are also interesting for some applications. Weak form efficiency merely requires that future stock returns may not be predicted from prior returns. The strong form of efficiency requires that stock prices reflect all available information, public and private. See Peterson (1994, 51-52).
7. Langbein and Posner (1980) discuss socially responsible investment in the context of efficient markets and the capital asset pricing model (CAPM). The CAPM provides a specific mechanism for adjusting stock returns to take account of risk. However, the efficient markets hypothesis does not depend on the CAPM or any other specific model of securities' returns.
8. Fama (1970, 1991) surveys the literature on market efficiency.
9. The author is unaware of any empirical studies supporting or refuting this claim.
10. Bruyn (1987, 13) provides a number of examples in which investing based on socially responsible criteria would yield superior results and in which socially responsible investment mutual funds outperformed the overall stock market over selected time periods. Hylton (1992, 28-32) lists the performance of a number of socially responsible investment funds and notes that some of them have outperformed both the S&P 500 index and the average of all mutual funds over one- and five-year horizons ending in the early 1990s. She also notes that some funds did not outperform the other indexes. See also Stoval (1992) and Kinder (1993) for some evidence on socially responsible investment funds outperforming other mutual funds.
11. For example, Galen (1994) notes that over a three-month period in 1994 a fund for "sinners" outperformed both a fund for "saints" and the average of all mutual funds. The sinners fund, Morgan Funshares, invested primarily in companies that specialize in industries such as alcohol, tobacco, and gambling. See also Teper (1991) for statistics suggesting that using social criteria reduces a fund's performance.
12. See chapter 7 in Peterson (1994) for a discussion of risk and return in finance.
13. The CAPM is a widely used model of stock returns. The primary insight of the model is that investors should earn a higher rate of return only for risks they cannot diversify away (nondiversifiable risk). See Grinblatt and Titman (1994) for an analysis of alternative methods of analyzing mutual fund performance.
14. An assumption maintained throughout the analysis is that corporations' management and boards of directors follow policies designed to maximize shareholder wealth. If they do not follow shareholder wealth maximization, then shareholders may be able to use their voting power to force changes in policies and thereby increase the firm's rate of return. Shareholder initiatives to change firms' management policies fall within the broad area of corporate governance, a topic of considerable academic research interest in recent years. Some of this research suggests substantial opportunities to improve many corporations' performance; however, these studies have generally focused on issues unrelated to that of following socially responsible investment principles. This study will not address corporate governance issues for two reasons: (1) little academic evidence exists to support (or contradict) the hypothesis that firms following socially responsible investment principals have superior earnings, and (2) the issues raised by corporate governance questions go far beyond the scope of this article.
15. An alternative is to calculate an internal rate of return (IRR) for the project and compare that with the firm's required rate of return. NPV has several theoretical advantages over IRR, but many businesses nevertheless use IRR. The choice of NPV or IRR for analyzing projects is unimportant for the purposes of this study. See chapter 9 of Peterson (1994) for a discussion of the use of NPV and IRR for evaluating projects.
16. This discussion assumes that the investment project is normal in the sense that it requires a large initial investment followed by a stream of cash inflows to the firm in future years. The analysis would change if the project were to involve a sufficiently large cash outflow at the end—for example, a substantial expense at the end of the project to

restore the environment to its original condition. In this case a large reduction in the value of future cash flows (a large discount rate) could actually encourage investment. Thus, an increase in the real rate of interest (market interest rate minus inflation) will have the effect of encouraging environmentally "dirty" projects that require clean-up in the future over projects that are environmentally "clean" from the start.

17. Elasticity of demand (supply) refers to the change in quantity demanded (supplied) in response to a change in the market price.
18. This statement is based on the currently plausible assumption that most investors (or at least the marginal investors) care only about the distribution of returns. If all investors incorporated both return distribution and a similar set of social values in their pricing, then "economic values" would depend at least in part on investors' social values. Note, however, that merely having a majority of investors follow socially responsible investment may not be sufficient for economic values to incorporate social values. Non-socially responsible investors may offset the socially responsible investors by holding less of the socially desirable stocks and more of the undesirable stocks. Non-socially responsible investors would follow such a strategy if it promised a higher risk-adjusted rate of return than holding stocks in proportion to their market value (which in finance terms is referred to as holding the market portfolio).
19. Indeed, a substantial body of academic finance literature has arisen in recent years under the heading of "market microstructure" to address the issue of how stock prices respond to changes in the order flow. See Kyle (1989) for an example of a model in which prices are not set in perfectly elastic markets. For a general survey of the market microstructure literature see O'Hara (1995) and the references therein.
20. Bond investors may be able to lower a firm's WACC by buying primary issues of the firm at below-market yields.
21. Indeed, the very large investment funds often avoid investment in small capitalization stocks because the large funds cannot buy or sell these stocks in sufficient quantity to influence their overall portfolio returns without causing large, albeit often temporary, changes in the small firm's stock price.
22. Although socially responsible investors may be able to bid up the price of a stock in the short run, in order to maintain the higher price socially responsible investors may need to purchase virtually all of the stock that reaches the market in the future. Investors who do not follow socially responsible investing are likely to be reluctant to purchase a socially responsible stock after its price has increased because they would be paying more for the stock than is indicated by its fundamental value. (Equivalently, the expected rate of return on the stock may decline after its price has been bid up, as is shown in the next section.)
23. The effectiveness of investors selling their shares may be enhanced by coordinated efforts to have consumers boycott the firm's products.
24. Scholes (1972) examines the impact of secondary stock offerings (large stock offerings that must be reported to the Securities Exchange Commission) for New York Stock Exchange (NYSE) stocks and finds that these offerings reduced share price but the magnitude of the decline was very small. More recently, Chan and Lakonishok (1993) examine the impact of institutional trades on NYSE and American Stock Exchange (AMEX) stock returns. They find a larger impact than Scholes and a bigger impact for buys than for sells. However, they still find that these often very large trades caused stock prices to move by less than 1 percent, even for the smallest firms in their sample.
25. Munnell, Blais, and Keefe (1983) make essentially the same point with regard to state investment in housing finance programs. Any state-sponsored program that effectively increases the supply of funds available to housing must result in the state fund earning a below-market rate of return.
26. See Starobin (1993) for a discussion of the politics of imposing CRA-type requirements on nonbank intermediaries.

References

Bruyn, Severyn T. *The Field of Social Investment.* New York: Cambridge University Press, 1987.

Chan, Louis K.C., and Josef Lakonishok. "Institutional Trades and Intraday Stock Price Behavior." *Journal of Financial Economics* 33 (April 1993): 173-99.

Fama, Eugene F. "Efficient Capital Markets: A Review of Theory and Empirical Work." *Journal of Finance* 25 (May 1970): 383-417.

_____. "Efficient Capital Markets: II." *Journal of Finance* 46 (December 1991): 1575-1617.

Fierman, Jaclyn. "The Coming Investor Revolt." *Fortune* 129, October 31, 1994, 66-76.

Galen, Michele. "Sin Does a Number on Saintliness." *Business Week*, December 26, 1994, 8.

Grinblatt, Mark, and Sheridan Titman. "A Study of Monthly Mutual Fund Returns and Performance Evaluation Techniques." *Journal of Financial and Quantitative Analysis* 29 (September 1994): 419-44.

Hamilton, Sally, Hoje Jo, and Meir Statman. "Doing Well While Doing Good? The Investment Performance of Socially Responsible Mutual Funds." *Financial Analysts Journal* (November/December 1993): 62-66.

Hylton, Maria O'Brien. "'Socially Responsible' Investing: Doing Good versus Doing Well in an Inefficient Market." *American University Law Review* 42 (1992): 1-52.

Kaufman, George G., and Larry R. Mote. "Is Banking a Declining Industry? A Historical Perspective." *Economic Perspectives* 18 (May/June 1994): 2-21.

Kinder, Peter D. "Social Investing's Strength Lies in Readiness to Deal with World's Tough Questions." *Pension World* 29 (April 1993): 10-12.

Knoll, Michael S. "Socially Responsible Investment and Modern Financial Markets." University of Southern California Law Center Working Paper 94-13, December 1994.

Kyle, Albert S. "Informed Speculation with Imperfect Competition." *Review of Economic Studies* 56 (July 1989): 317-56.

Langbein, John H., and Richard A. Posner. "Social Investing and the Law of Trusts." *Michigan Law Review* 79 (November 1980): 72-112.

Munnell, Alicia H., Lynn E. Blais, and Kristine M. Keefe. "The Pitfalls of Social Investing: The Case of Public Pensions and Housing." Federal Reserve Bank of Boston *New England Economic Review* (September/October 1983): 20-37.

O'Hara, Maureen. *Market Microstructure Theory.* Cambridge, Mass.: Basil Blackwell, Inc., 1995.

Peterson, Pamela P. *Financial Management and Analysis.* New York: McGraw-Hill, Inc., 1994.

Scholes, Myron S. "The Market for Securities: Substitution versus Price Pressure and the Effects of Information on Share Prices." *Journal of Business* 45 (April 1972): 179-211.

Srinivasan, Aruna. "Intervention in Credit Markets and Development Lending." Federal Reserve Bank of Atlanta *Economic Review* 79 (May/June 1994): 13-27.

Starobin, Paul. "Make 'Em Pay." *National Journal,* July 24, 1993, 1856-61.

Stoval, Robert H. "When Do-Gooders Do Good." *Financial World* 161, September 1, 1992, 68, 70.

Teper, Jeffrey A. "The Cost of Social Criteria." *Pensions and Investments,* May 13, 1991, 34.

Tschinkel, Sheila L., and Larry D. Wall. "Some Lessons from Finance for State and Local Government Development Programs." Federal Reserve Bank of Atlanta *Economic Review* 79 (January/February 1994): 1-10.

Inflation and Inflation Forecasting: An Introduction

Ellis W. Tallman

Inflation in the United States has been moderate and relatively stable for the past several years, hovering at an average annual rate of around 3 percent. From the perspective of recent history, the currently modest rates of inflation are a vast improvement over the high and occasionally double-digit inflation rates of the 1970s and early 1980s. Since 1973 the inflation rate, measured as the rate of change from the year-ago quarterly Consumer Price Index (CPI), has averaged approximately 6 percent.

Despite the relatively benign rates of inflation currently observed, inflation and the fear of inflation continue to penetrate the business world. Almost weekly, major movements in the stock and bond markets are attributed to economic data announcements that allegedly influence market perception of the future inflation rate. Clearly, the lower rate of inflation observed recently does not greatly reduce the need or desire for inflation forecasts; the costs of unexpected inflation on individual firms and consumers remain substantial even at low levels of inflation.

Expectations about inflation are embedded in planning decisions of all kinds. Labor unions need inflation forecasts to help refine their wage demands, and firms planning for future expenses and expected revenues need

The author is an economist in the macropolicy section of the Atlanta Fed's research department. *He wishes to thank B. Frank King, Eric Leeper, and Mary Rosenbaum for comments. The author is solely responsible for contents.*

to account for expected inflation. Government budget-making has perhaps an even greater need for accurate forecasts of inflation, given that Social Security payments are linked directly with the inflation rate measured by the CPI. Consider, for example, the implications of the following scenario. In 1993 the aggregate level of Social Security expenditures was approximately $305 billion. Suppose the government had budgeted Social Security expenditures for the next ten years to increase at 3 percent per year but, because of higher than expected inflation, the actual increase was 5 percent. The cumulative difference in the amount of Social Security payments anticipated and the amount paid in these ten years would be nearly $400 billion.[1] Recent attempts at legislation to enforce a balanced federal budget only heighten the importance of accurate inflation forecasts for planning expenditures and revenues.

This article provides an overview of the effects of inflation and the significance of inflation forecasting, first considering some of the reasons that businesses may need to forecast inflation. The goal is to provide a "consumers guide" to forecast accuracy and evaluation. The discussion describes a selection of forecasting methods, some of which use both statistical models and forecaster judgment, or "skill." The analysis examines as examples two statistical models that employ standard economic theory to suggest variables that help predict inflation and reviews the subsequent inflation forecasting models. Forecasts from each model are then compared in a variety of ways with simple univariate time series models that use only past data on inflation to forecast the future without reference to economic theory.

The Uses of Inflation Forecasts

The design of inflation forecasts reflects the varying needs and concerns of businesses, financial market participants, government policymakers, and others who are the end users of those predictions. Purely statistical models (namely, time series models) employ quantitative techniques to extract information from observations of past data in order to generate a forecast without the use of economic theory. Economic models use economic theory to justify the choice of explanatory variables and help explain the behavior among the variables that underlie the forecast. In designing a formal forecasting model, the forecaster considers the user's needs in making choices about several issues: (1) the time horizon over which the model must forecast (one month, one quarter, one year, and so forth), (2) the data frequency for the forecast (for example, monthly, quarterly, or annual data), (3) the type of model (a time series model or a model that employs variables that hypothetically cause inflation), (4) the method for evaluating forecast "accuracy" (for example, the root mean squared error of the forecast), the forecast bias (whether on average the forecasts over- or underpredict inflation), or the probability of predicting a turning point (say, from low to high inflation).[2]

Models are specifically designed to fill the needs of particular forecast users. For example, a private business like a manufacturing firm requires a significantly different kind of inflation forecast than a policymaker or a financial market participant. Such businesses typically use forecasts to help monitor input costs, revenues from output sales, the real (inflation-adjusted) level of profits, and the like. Because businesses generally employ forecasts with a medium to long-term time horizon—from two to five or ten years out—quarterly or annual data are likely to be satisfactory for their planning purposes. An additional consideration is that businesses rarely view inflation as an end result of their decisions or behaviors so that their needs may be adequately met by purely statistical models.[3] The primary concern of a business inflation forecast will be forecast accuracy, perhaps as measured by the typical evaluation criteria, but it may be a combination of a number of them. In general, models for businesses use the root mean squared error criterion because the criterion penalizes large errors and thus helps avoid big misses in their forecasts.[4]

A financial institution has quite different inflation forecasting needs. For instance, it may want to predict the probable future value of its fixed-rate loan portfolio. Similarly, financial portfolio managers use inflation forecasts to determine the best asset allocation for maximizing real asset returns, perhaps substituting out of assets whose real returns erode amid inflation if the forecast signals increasing inflation. The desired forecast horizons may range from the short (monthly) to the very long term (thirty years) for these forecast consumers, and models designed for them are likely to employ most of the available data frequencies. Both interpretable macroeconomic and purely statistical models are used.

Inflation forecasts are also important to bond traders, who trade in markets that react quickly to new information. For a forecasting model to be useful in their attempts to make profits, bond traders may need to know only the direction that inflation will move (relative to market anticipations) in the next inflation rate announcement. For example, the news services survey market participants' forecasts prior to an inflation rate announcement and then publish the survey average as the market forecast. Traders have an interest in a model that can improve on the average forecast. For such a short forecasting horizon, the models will use monthly data. The fact that only a limited number of macroeconomic data are published monthly perhaps constrains inflation prediction models that use other macroeconomic variables.

Because accuracy of inflation forecasts is an important consideration for bond trading operations, a number of forecasting techniques are likely to be used in models designed to meet their needs. As in models for private businesses, the root mean square error may be the key measure of the accuracy of the bond trader forecast because it helps reduce the likelihood of big errors. Judgment is another key component of their forecasts because of the absence of timely data. Any descriptive short-term data relationships—that is, relationships that appear in the data but that may have no deeper economic meaning—may help predict the next month's inflation rate, and there is no need for any interpretability in the model.

Monetary and fiscal policymaking institutions, working without the profit motive of private businesses and with social goals, have a different objective in using inflation forecasts. One motivation for a policymaker is to attempt to minimize the social costs imposed by inflation. Economic models chosen for analysis of inflation as well as for inflation forecasting typically take into account some role for policy variables in the inflation process. Policymakers typically assume that policy choices affect inflation's behavior, and a model of inflation that will be adequate for their purposes is therefore likely to provide some insight into policy's potential impact on inflation.[5] The artificial model of the inflation process then becomes the basis of the inflation forecasting model. The time horizon for an inflation forecast to serve policymakers should be long enough to capture the effects of policy on the inflation process, typically assumed to extend as far as two to five years.

In most cases, the final inflation forecast numbers for any of the entities mentioned above include adjustments to model forecasts that apply the judgment of the forecasters, that is, changes initiated to account for information not contained in the economic or purely statistical model. Including such judgmental adjustments in forecast numbers makes distinguishing among forecasting models ambiguous because it is difficult to correct for the forecaster's judgment versus the model implications.[6]

Macroeconomic Inflation Forecasting Models

Of the two main types of inflation forecasting models—purely statistical models and economic (macroeconomic) models—economic models have some characteristics that make them typically more desirable. As discussed, purely statistical models employ past inflation data to forecast future inflation, and there is no additional explanation of the forecast.[7] Economic models, on the other hand, are designed to be interpretable so that movements in the explanatory variables explain the inflation forecast. Also, inflation forecasts that use data in addition to inflation data contain numerous correlations that purely statistical models cannot.[8]

This article focuses on the forecasting performance of two simple economic approaches in describing the inflation process. Using simple formulations allows emphasizing the intuitive appeal of each model and highlighting the contrast between them. While adding other variables to the specification may improve the models' forecasting performance, doing so would hinder the direct inferences about the contribution of the intuition of the model to forecasting performance.

Two representative macroeconomic paradigms typically underlie common empirical models of inflation and appear comparably successful at forecasting inflation.[9] These two standard models incorporate interactions among variables that are hypothesized to influence inflation significantly.[10] This discussion emphasizes the main intuitions from the theoretical models as they translate into the model specifications typically estimated for forecasting inflation.[11] Simple models of inflation are examined with a focus on what the key elements of each approach contribute to forecasting inflation.

One popular macroeconomic model used for forecasting inflation is the so-called Phillips curve.[12] The original Phillips curve relation observed the negative correlation of the rate of change in money wages in the United Kingdom with the country's unemployment rate. Over time, this empirical association has been similarly applied to aggregate inflation and aggregate output.

The concepts of "potential output" (or potential gross domestic product) and the "natural rate of unemployment" help describe the aggregate application of the

Phillips curve. The potential output measure indicates the hypothesized amount of output that the economy could generate if it produced at full capacity; it is typically referred to as the full-employment level of output. An alternative description employs what is called the natural rate of unemployment—that is, the rate of unemployment that persists in periods of healthy economic growth, such as unemployment resulting from poor worker-job matches that result in resignations or temporary frictional unemployment and from inadequate worker training relative to existing job opportunities, known as structural unemployment. The Phillips curve relationship in an aggregate formulation suggests that an unemployment rate lower than the natural rate of unemployment or a growth rate of real GDP that surpasses that of potential GDP is associated with a higher inflation rate. The concepts of potential output and the natural rate of unemployment are nearly interchangeable in the Phillips curve inflation model.

More recent formulations of the Phillips curve approach suggest that the underlying potential rate of real growth represents the highest rate at which the economy could grow without increasing inflation. In this view, output growth at this potential rate does not exert inflationary pressure. If the economy grows at a faster rate, firms face production capacity constraints that force them to hire additional labor and to work capital more intensively. In this process, firms bid up the price paid to production factors (namely, wages), and production costs increase. Firms pass the supply price increases through to the final goods prices consumers face, causing inflationary pressures to increase.

The basic idea behind the Phillips curve is that if demand causes output to exceed the measure of potential output, the economy gets stretched beyond its capacity for noninflationary growth. The prices of then-scarce production factors are bid up, and the result is increased prices that lead to reduced output demand. In contrast, if output growth is below potential, inflationary pressures are reduced, and inflation should decline. The difference between the real economy's potential output and actual output is referred to as the "gap" or, more specifically, the GDP gap.

Empirical models that apply the Phillips curve intuition typically emphasize the central role of the GDP gap in the inflationary process.[13] The empirical tests in the research reported on here allow in the prediction equation only lagged values of all variables that in the initial specification are contemporaneously determined along with inflation. The number of inflation lags is limited to four.[14] The estimated relationship is

$$\Delta p_t = \alpha_0 + \beta_1(y_{t-1} - y_{pt-1}) + \beta_2(\Delta Y_{t-1} - \Delta y_{pt-1}) + \Sigma_{j=1}^{4} \beta_2 \Delta p_{t-j} + v_t.$$

The Phillips curve inflation forecasting models often estimate single-equation models like the one above, suggesting that these simple models capture the key relationships in the data that determine inflation. But single-equation models are inherently partial equilibrium models, which hold many other variables constant by assumption, failing to account explicitly for monetary or fiscal policies and their potential effects on inflation. Also, the model is essentially a demand-based model; potential output reflects the level of output that the economy can supply without inflation, and demand conditions determine whether inflation increases or decreases.

The second macroeconomic model presents an alternative view of the inflation process. The traditional monetarist approach relies on the intuition that observations of past growth rate in the money supply predict the long-run inflation rate. Monetary inflation models suggest the following simple relationship:

$$\Delta p_t = f(M_t, \Delta M_{t-1} \ldots M_{t-i}),$$

where Δp_t is the inflation rate for period t and ΔM_t is the money growth rate for period t. The simple relationship suggests that growth in the aggregate supply of nominal money determines the inflation rate. Inflation is therefore some function of current and past money growth measures.

Most empirical applications of the monetarist intuition suggest that the relationship between money growth and subsequent inflation is a single linear equation: the past and current growth rates of money translate directly into subsequent inflation.[15] Empirical estimations (and forecasts) of monetarist forecasting models depend largely on the choice of monetary aggregate used to measure money growth, typically M1, M2, or the monetary base. The appropriate monetary measure (monetary aggregate) as well as the number of lagged observations of money growth for the inflation equation (what constitutes the "long run") are not suggested by theory but significantly affect the forecast performance of the model.

The results in Yash P. Mehra (1988), David J. Stockton and James E. Glassman (1987), and William Reichenstein and J. Walter Elliot (1987) suggest that this simple monetarist formulation forecasts relatively poorly regardless of the choice of monetary aggregate. Rather than use such a formulation, the present research employs an inflation forecasting model inspired by a simple money demand relationship presented in Eugene F. Fama (1982) and in a more extensive form in Elliot and Reichenstein (1987). The modeling strategy bears a direct resemblance to the standard monetarist model because of the appearance of a monetary aggregate in the specification.[16] However, the framework includes several other variables that help explain the inflation process. For this article, the demand for real balances (the money supply deflated by the price level) is

$$M^d/P = f(y, i).$$

The demand for real money balances increases as the level of real activity increases, that is, as output grows. Conversely, increases in the rate of interest raise the opportunity cost of holding money balances (assumed to be non-interest-bearing) instead of interest-bearing assets and thus lower the demand for real balances. This simple relationship provides the underpinnings for another inflation forecasting model that uses a single linear equation, the money demand inflation forecasting equation.

The intuition underlying the money demand inflation forecasting equation is that a rate of money supply growth that is the same as the rate of growth determined by interest rate movements and output (money demand) will not increase inflation. When the monetary aggregate growth rate exceeds the rate at which money demand grows (determined by the interest rate and output movements), the excess money growth affects the price level variable and results in inflation.[17] The money demand model describes a simple mechanism of inflation generation but does not attempt to uncover variable interactions that underlie the general intuitive story.

Empirical specifications of the model vary, depending on the forecasting problem at hand. For the purposes of this study, the estimated money demand model has four lags of the explanatory variables, meaning that the four most recent values of the explanatory variables are included in the estimated equation. The initial specification of the money demand inflation model has no lags of the dependent variable. By adding four lags of inflation to the right-hand side of the equation, the model becomes more comparable to both the Phillips curve model and the time series model specification (see below), each having in common four lags of inflation as explanatory variables.

Estimates of the single-equation money demand model of inflation suffer from criticisms similar to those aimed at the simple Phillips curve model, namely, that the model is partial equilibrium and cannot isolate the underlying sources of shocks that generate inflation in the economy. The framework is exclusively for a money demand model and cannot separately account for aggregate supply shocks like the oil price shocks observed in the 1970s. And even though a monetary aggregate measure is used in the model, the data measure does not capture changes in monetary policies, often referred to as "shocks." Neither can the simple model isolate the source of fiscal policy shocks that affect the movements in the explanatory variables. This inability of partial equilibrium models to isolate shocks severely limits the degree to which the models may be interpretable—that is, they are not structural. These single-equation models assume that the explanatory variable values are somehow known before inflation is forecast.

The macroeconomic model forecasts will be compared with those of purely statistical models. Univariate statistical models, needing no other variables than inflation in order to produce an inflation forecast, are easy to use. Results compared with economic model forecasts indicate whether either of the theory-based forecasting techniques improve significantly on the statistical model forecasts, which rely on past inflation data only. For simplicity, a simple autoregressive model of inflation with four lags of inflation is estimated:[18]

$$\Delta p_t = \alpha_0 + \textstyle\sum_{j=1}^{4} \beta_j \Delta p_{t-j} + \epsilon_t.$$

Model Estimation and Forecast Evaluation

All the single-equation models are estimated over the data sample from the fourth quarter of 1960 to the third quarter of 1994.[19] Certain shortcomings of the models can be detected from the full sample estimations. Specifically, the money demand specification without lagged values of inflation suffers from autocorrelated errors, that is, errors that are correlated with past errors, indicating a problem of persistent prediction errors.[20] Such a pattern in forecast errors suggests a serious problem in the forecast model design.

Some forecast evaluation procedures employed in the economics literature make strong, artificial assumptions about the availability of data on variables other than inflation. For example, in what is described as dynamic, out-of-sample forecasting, single-equation macroeconomic forecasting models may employ the actual values of the explanatory variables to generate the forecasts. In forecasts of one period into the future, this strategy is adequate if the data on all explanatory variables are available before the inflation measure. In real time—that is, forecasting today what will not be known until the future—using future values that could not be known at the time of the forecast is uninformative about the model's forecast accuracy. The unknown values of the explanatory variables should be forecast as well to generate a more realistic test of a model's forecasting accuracy.

The first forecasting exercise for the inflation models is to predict the inflation rate that will take place one quarter into the future. To investigate the relative accuracy of the models, out-of-sample forecasting is used in which the model is estimated over the sample period up until the first forecasting period, in this case, the fourth quarter of 1972.[21] Then, the model is forecast one period into the future, the first quarter of 1973. That forecast value is stored for later comparison with actual data. The estimation sample is then updated to include the first quarter of 1973, and the forecast process is repeated. After the model produces forecasts for all periods in the forecast sample, the series of forecast values can be compared with actual inflation data. Because the estimated models involve only lagged variables as predictor variables, the forecasts of each model can be compared with a real-time forecasting problem in which inflation is forecast using only available data.[22] The results of the exercise are in Table 1.

Table 1
Forecast Evaluation One Quarter into the Future

	Mean Error	Root Mean Squared Error	Theil U
Forecast Sample: 1973:1-1983:4			
Phillips Curve	.426	2.263	.889
Money Demand Plus	–.147	2.135	.839
Time Series	.179	2.413	.948
Money Demand	.874	2.478	.974
Forecast Sample: 1984:1-1994:3			
Phillips Curve	.097	1.603	.938
Money Demand Plus	–.537	1.788	1.026
Time Series	–.235	1.595	.933
Money Demand	–1.940	2.850	1.666
Full Sample: 1973:1-1994:3			
Phillips Curve	.275	1.983	.907
Money Demand Plus	–.279	1.980	.905
Time Series	–.017	2.065	.944
Money Demand	–.473	2.670	1.220

The evaluation criteria compare the forecast values with the actual value of inflation and generate summary measures of forecast accuracy. As mentioned above, the root mean squared error is often the key measure of forecasting accuracy. However, one can also use the mean error to gauge whether the forecasting model generally over- or underpredicts the inflation rate. Depending on how the forecaster perceives the cost of errors, the mean absolute error may be preferred to the root mean squared error. The mean absolute error does not weigh a large error as heavily as does the root mean squared error (which imposes a quadratic cost of error function).

Although all evaluation statistics are listed for the forecasts, the discussion will focus on a measure of relative accuracy, the Theil U statistic. The Theil U statistic represents the ratio of the root mean squared error of the given forecasting model to the root mean squared error of a naive "no-change in the dependent variable" (meaning, "the same inflation rate as observed last time") model, one that an individual naive about economic theory can employ as a forecasting mechanism. As a criterion for a useful model, one would hope that an inflation forecasting model can outperform this naive forecast. If the Theil U statistic is greater than one, then the naive model outperforms the given forecasting model. Values of the Theil U statistic close to one suggest that the model under scrutiny performs no better than a no-change forecast. On the other hand, values of the Theil U statistic less than one imply that the forecasting model forecasts more accurately on average than the naive model. The article focuses on the Theil U statistic as a way to summarize and compare root mean squared error, the chosen criterion to distinguish among rival models.[23]

Table 1 shows that the Phillips curve model and the money demand plus inflation lags model perform similarly over the forecast subperiod from the first quarter of 1973 to the fourth quarter of 1983 as well as over the entire forecast sample. For the full sample, the Theil U statistics are similar and less than one, suggesting improvement over the simple "no-change" forecast and slight improvement over the simple time series model. Also, the Phillips curve model has a mean error that is positive (underpredicting inflation on average), whereas the mean error of the money demand model is of similar size but of opposite sign (overpredicting). In contrast, the mean error of the time series model is close to zero for the full forecast sample.

For the latter sample (the first quarter of 1984 to the third quarter of 1994), the Phillips curve model has a lower Theil U statistic than the money demand and a smaller mean error that is close to one. Over this subsample, the time series model also appears to outperform the money demand model, having a lower mean error and a lower Theil U statistic. In contrast, the Theil U statistic of 1.666 for the money demand model without inflation lags suggests that the model forecasts significantly worse than a no-change forecast in the latter forecast sample.[24] The mean error of –1.940 indicates that the money demand model (without inflation lags) is poorly specified, and it is not examined further.

The next exercise for the models is to forecast four quarters into the future. The longer forecast horizon requires the use of data that in a real-time forecast would be unavailable. In this procedure the model is estimated over the sample period, forecast four quarters into the future, and those forecast numbers are accumulated into one forecast of the average inflation rate expected to persist over the next year. The process is iterated as above to create a series of one-year ahead forecast numbers. It is important to note, however, that the number of nonoverlapping data points or "observations" is much smaller (the forecast sample divided by four, or twenty-two yearly observations [if 1994 is included]) than in the one-step ahead forecast exercise.[25]

There are two sets of results for the model evaluation statistics. The first set is evaluation statistics from a static forecast that assumes that the actual values of data in the forecasting model are known (essentially, getting new information on the actual explanatory variables as the prediction extends further into the future, but information that is unavailable for forecasting in a real-time setting). The second set of results allows partial dynamics—that is, all the models have lagged inflation as an explanatory variable so that the model forecasts of inflation are used as predictors, but all other explanatory variables employ actual data values in the prediction equation. A fully dynamic model generates forecast values for future observations of the explanatory variables using the interrelationships among the data to forecast them. This approach contrasts with partially dynamic models that use future actual values, which are obviously unavailable in a real-time forecast exercise. Forecast evaluation statistics are listed in Table 2.[26]

The most noticeable feature of the data in Table 2 is how much smaller the forecast error statistics are for the static forecasts than for the partially dynamic forecasts. The Theil U statistics for the static forecasts are

Table 2
Forecast Evaluation One Year into the Future

	Mean Error		Root Mean Squared Error		Theil U	
	Static	Partially Dynamic	Static	Partially Dynamic	Static	Partially Dynamic
Forecast Sample: 1973:1-1994:3						
Phillips Curve	.277	.449	1.053	1.943	.463	.855
Money Demand Plus	–.335	–.584	1.089	1.923	.480	.847
Time Series	–.063	–.137	1.038	2.256	.457	.993
Forecast Sample: 1984:1-1994:3						
Phillips Curve	.049	.080	.683	1.292	.497	.940
Money Demand Plus	–.447	–.737	.942	1.649	.685	1.199
Time Series	–.254	–.514	.676	1.472	.491	1.071

almost half those of the partially dynamic forecast Theil U statistics, suggesting that the restriction of employing a forecast of inflation rather than future actual values of inflation almost doubles the variance of the forecast. Another point worth noting is that the mean errors are uniformly smaller in the static forecast exercise than in the partially dynamic forecasts, indicating that the addition of some dynamics allows the models to stray more often.

In the partially dynamic results, both the Phillips curve and money demand with inflation lags have Theil U statistics less than one, clearly better performance than the no-change forecast and somewhat better than the time series model. For the full sample results, the mean errors of the respective models are of opposite signs, 0.449 for the Phillips curve and –0.584 for the money demand models, similar to the one-quarter-out forecasts. Charts 1 and 2 present the forecast errors for the money demand and Phillips curve models, respectively, over the full forecast sample. Positive forecast errors for both models appear largest around the two oil shocks (1973-74 and 1979-80), indicating underprediction of inflation and that these simple demand models cannot adequately account for supply shocks. Although both models overpredict inflation during the 1975-76 recession and in the 1981-83 recession, the money demand model appears to have persistent overpredictions of inflation throughout the early 1980s, a period in which many observers noted a dramatic shift in monetary policy. Some observers viewed the change in Federal Reserve policy (and possibly policy goals) in 1979 as affecting the structure of the relationships among the data.[27] The single-equation money demand model appears unable to account for the change in monetary policy and the resulting effects on inflation.

In the shorter forecast sample from 1984:1 to 1994:3, the money demand specification has a Theil U statistic of approximately 1.2, suggesting that the model forecasts less accurately than the "no-change" forecast. Also, the mean error of –0.737 indicates that the money demand model overpredicts inflation on average, which is undesirable if the forecast objective is an unbiased forecast. Over this later sample, the Phillips curve model appears to forecast more accurately than the alternative models and the "no-change" forecast.

As mentioned above, there are unrealistic assumptions about data availability that are implicit in these results. The use of actual data for the future observations of the explanatory variables in the forecasting equation offers an unfair advantage to the macroeconomic forecasting models relative to the time series

Chart 1
Inflation Forecast Errors for Money Demand Models
(*Static, Partially Dynamic, and Dynamic Forecasts*)

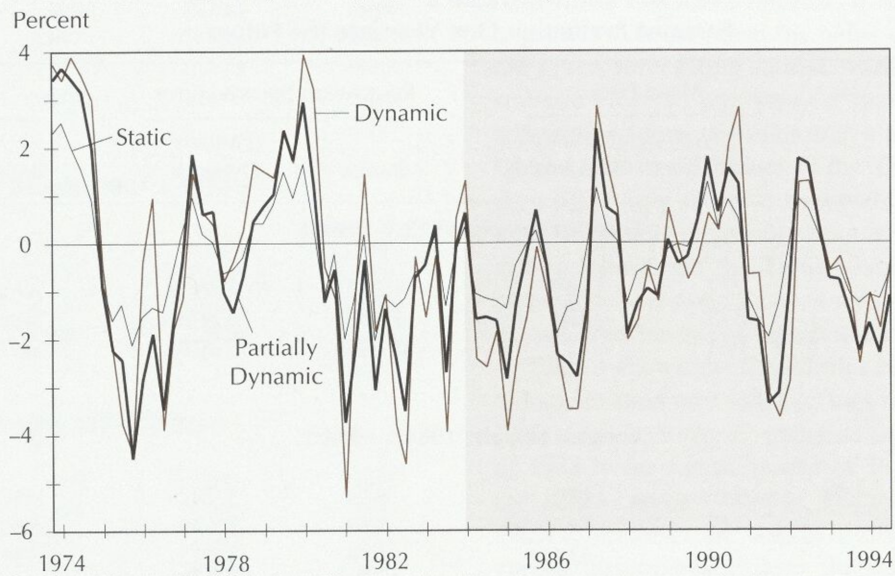

model. Given that advantage, the poor performance of the money demand model in the latter sample suggests even weaker forecast results when the explanatory variables are forecast. The above results suggest that the relationships underlying the money demand model may have weakened or changed in the latter sample, perhaps as a result of the monetary policy shift, financial innovation, or both.[28]

*R*epresentative Dynamic Models of Inflation

The simple linear forecasting models of both macroeconomic approaches fail to specify how to generate values of the right-hand-side variables when the model forecasts out-of-sample, that is, when there are no data available, as in a forecast two or three periods into the future. A more realistic procedure is to use an explicitly dynamic model to forecast inflation out-of-sample. Unfortunately, the simple models have no suggestions about how to specify forecasting models for the explanatory variables in the model. Stockton and Glassman (1987) employed time series specifications to forecast values of the explanatory variables in their comparison of inflation forecasting models. For this article, forecasts of the Phillips curve model were generated by specifying simple equations for the growth in nominal GDP and the level of real GDP and then imposing identities to create the GDP gap ($y_{t-1} - y_{pt-1}$) and the nominal growth gap ($\Delta Y_{t-1} - \Delta y_{pt-1}$) variables.[29] For the money demand model, a vector autoregression (VAR) is estimated that includes the money demand plus inflation lags as one equation of the dynamic system. The VAR model generates inflation forecasts in a dynamic setting that more closely mimics the real-time forecasting problem.[30]

In a vector autoregression, every variable on the right-hand side has a prediction equation associated with it. Thus, by design the model performs a truly dynamic forecast of inflation out-of-sample because it forecasts all variables in the system.[31] VAR models forecast in a pseudo-real time setting. That is, they forecast values of the explanatory variables implicitly in a dynamic forecast of inflation. The VAR generates forecast values of the variables in the inflation equation, and the accuracy of the inflation forecasts relies on how accurate the VAR forecasts the other variables in the system.[32] The forecast evaluations of the VAR present more realistic forecast accuracy statistics for the dynamic forecasts of the average inflation rate for

Chart 2
Inflation Forecast Errors for Phillips Curve Models
(*Static, Partially Dynamic, and Dynamic Forecasts*)

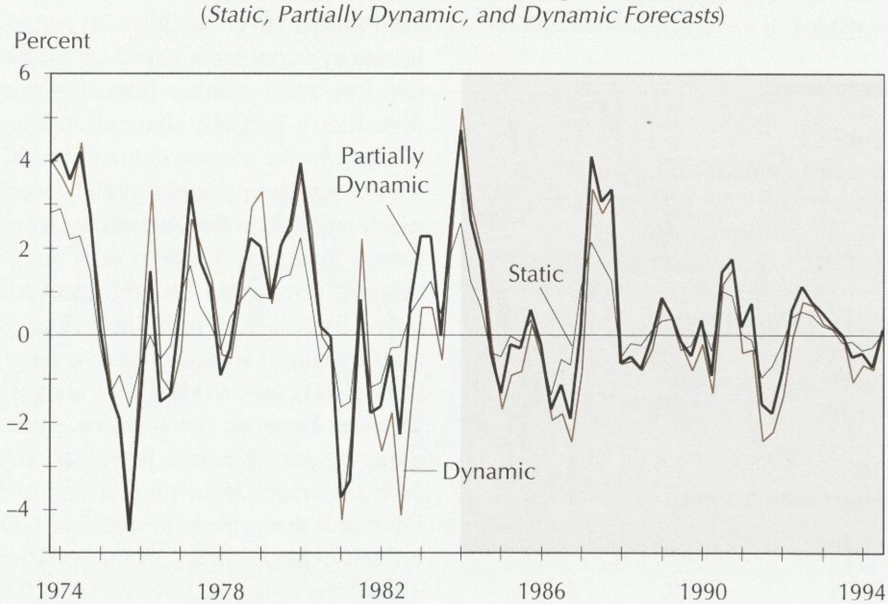

next year. The money demand model with inflation lags is essentially the equation that forecasts inflation from a vector autoregression model that contains all the variables in the forecasting equation (that is, interest rates, output growth, monetary aggregate growth, and inflation). These results appear in Table 3.

The VAR dynamic forecast fails to improve significantly over the naive no-change forecast in the full sample results. Also, the root mean squared error of the VAR is larger than the root mean squared error for the single-equation version that employs actual variables as explanatory variables. Thus, the increase in the root mean squared error measure can be attributed to the forecast errors from forecasting the explanatory variables. It can be seen more clearly how the forecast errors for inflation increase as the realism of the forecasts in Chart 1 is increased. Note that this chart refers to the money demand specification.

Chart 1 presents the forecast errors from static forecasts (forecasts using only actual values), partially dynamic forecasts (forecasts that allow only inflation forecasts in forecasting further in the forecast horizon), and dynamic forecasts (forecasts using forecast values of all variables in the equation). The plots highlight three noticeable features of the forecast error series. First, although the static forecast errors are clearly the smallest, the static forecast still makes sizable errors, noting that the model errs despite the benefit of forecasting with actual future data not available in a real-time forecasting exercise. Secondly, the partially dynamic forecast errors are substantially larger than the static forecast errors, suggesting that the use of inflation forecast values in subsequent forecasts of inflation compounds the initial forecast errors and worsens the forecast performance. Finally, the dynamic forecast errors are larger than the two others, especially when the forecast errors are large, emphasizing the contribution of the explanatory variable forecast errors to the forecast error of the inflation forecast.

The shaded area of Chart 1 highlights the forecast errors of the three forecasts over the sample period 1984:1 to 1994:3. Over this sample, the forecast errors for the dynamic model are generally larger than the errors of the other two models. Unfortunately, only the dynamic model forecast presents a reasonably realistic test of forecasting accuracy, and over this sample the VAR forecasts relatively poorly. This example illustrates one cost of additional variables in a forecasting equation—that is, those data series must be forecast into the future as well in order to forecast inflation. An explanatory variable that is virtually unforecastable may help the statistical fit in-sample but actually hurt the out-of-sample forecasts.[33] Over the 1984:1-1994:3 forecast sample the VAR performs substantially worse than the no-change forecast, overpredicting inflation by nearly 4 percent during the 1990-91 recession.

For the dynamic forecast of the Phillips curve model, the Phillips curve inflation equation is estimated as well as regression specifications for the log level of real GDP and the growth rate of nominal GDP, respectively, in equations that do not differ greatly from a VAR. From these equations, identities are employed to generate the GDP gap and the nominal GDP growth gap.[34]

The full sample forecasting results for the dynamic Phillips curve model are only slightly different from the Phillips curve model with actual data for the explanatory variables. Chart 2 shows how the static forecast has much smaller forecast errors than either the dynamic or partially dynamic forecasts. However, in contrast to the money demand model chart, this chart shows that the dynamic model forecast errors are not much larger than the forecast errors of the partially dynamic forecasts. These results suggest that inflation forecast errors that become incorporated into subsequent forecasts of inflation in the partially dynamic (and dynamic) forecasts have a large influence on inflation forecast errors. There are relatively small differences between the forecast errors of the partially dynamic and dynamic forecasts, suggesting that the forecast errors from predictions of the level of real GDP and the growth in nominal GDP do not greatly influence the accuracy of the inflation forecast. In addition, the dynamic forecast has a Theil U statistic of 0.896 for the full sample and of 0.980 for the latter half of the sample, comparable to the Theil U statistic

Table 3
Dynamic Forecast Evaluation

	Mean Error	Root Mean Squared Error	Theil U
Forecast Sample: 1973:1-1994:3			
Phillips Curve	.271	2.036	.896
Money Demand	–.492	2.220	.978
Forecast Sample: 1984:1-1994:3			
Phillips Curve	–.233	1.347	.980
Money Demand	–.805	1.880	1.368

of 0.855 and 0.940 for the partially dynamic forecast over those forecast samples. The shorter sample period has less variation in the inflation rate. As a result, the no-change forecast provides a relatively tougher standard for the other model forecasts to improve upon.

Chart 3 presents the dynamic forecasts of both the Phillips curve and the money demand models along with the actual one-year rate of inflation. In the latter part of the sample, the Phillips curve model more closely follows the path of inflation than does the money demand model forecast. Neither model appears to forecast accurately. The two models often miss the inflation rate in opposite directions in the early 1980s, the Phillips curve largely underpredicting the inflation rate following the serious recession of 1981-83 whereas the money demand model overpredicts inflation at that time. Later in the forecast sample, both models appear to err in the same direction.

The Theil U statistics reinforce the apparently better forecasting of the dynamic Phillips curve model over the VAR; the Theil U statistics of the dynamic Phillips curve forecasts are clearly lower than those of the VAR over the respective forecast samples.[35] Also, in the 1984:1-1994:3 sample period, the mean error of the Phillips curve model is much less than that of the VAR, suggesting a less biased forecast.

The statistical results suggest that the Phillips curve model may be a better forecasting model in real-time circumstances than the other models presented, including the simple time series model. However, the forecast evaluation statistics for the dynamic Phillips curve model were the best of the models examined but were not much better than the no-change forecast. For example, the Theil U statistic for the 1984:1-1994:3 forecast subsample was close to one, suggesting that the dynamic Phillips curve forecast was not a large improvement over the naive no-change forecast.

The VAR includes data series on all the variables in the Phillips curve model except for the potential GDP measure. Adding the potential GDP variable to the VAR makes the Phillips curve specification a restricted version of the more general VAR model. Forecasts of the VAR that include the potential GDP measure as a deterministic variable not only fail to improve but in fact worsen the forecast evaluation statistics for both forecast samples. This result reflects how imposing zero restrictions on a model (like the VAR) may produce improved forecasts. It is unclear, however, whether the resulting forecast (from the Phillips curve model) will be any more useful.

To evaluate further the effectiveness of a Phillips curve inflation forecasting model, the model needs to be tested in real-time forecasting exercises.[36] Thus, the

Chart 3
Inflation Forecast from Dynamic Models
(*Phillips Curve versus Money Demand*)

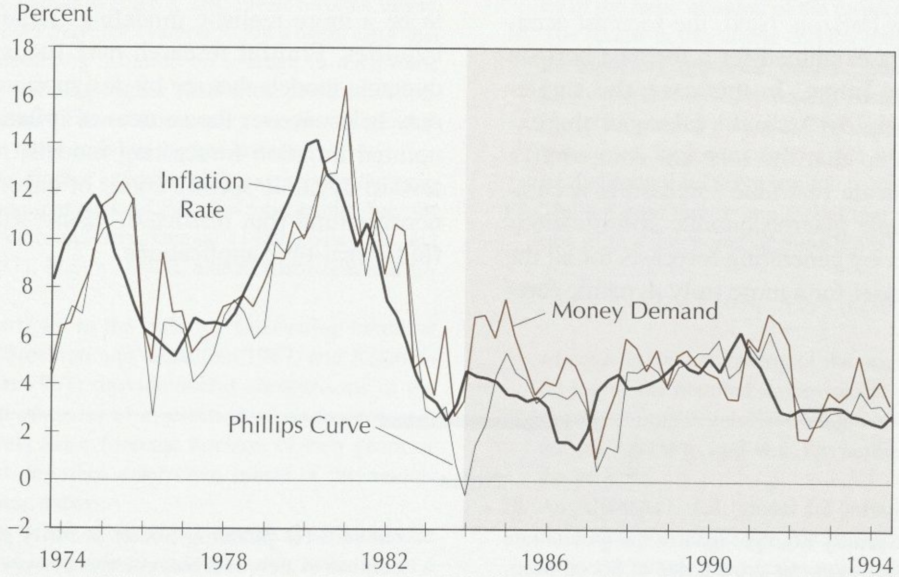

only true test of the model's forecasting ability is to use it in real-time forecasts and evaluate the results. Nevertheless, the statistical results in this article suggest that forecasting inflation accurately in real-time situations with simple economic statistical forecasting models is no easy task.

The results illustrate that the evaluation of forecast performance can change substantially by relaxing some of the unrealistic assumptions imposed in the forecast evaluation procedures. This research estimates the simple models imposing the unrealistic assumptions on data availability, that is, those using actual values on the predictor variables. Dynamic specifications of the two models, each of which generates a completely dynamic out-of-sample forecast, are then employed. Unfortunately, using the simple models here, the closer forecasting models move toward approximating the real-time forecasting exercise, the less impressive the empirical results for the causal forecast models.[37]

Conclusion

This article examines the forecast performance of two simple frameworks for modeling inflation—the Phillips curve and the money demand/monetarist models—in a variety of settings. The one-period (one-quarter) ahead forecasting accuracy of the models is examined first, and the results generally favor both macroeconomic models relative to the naive "no change in the inflation rate from the current inflation rate" in this forecast horizon. Next, the forecast accuracy of the models is examined for a forecast horizon of one year into the future. In this case, the single-equation models employ "actual" values of the explanatory variables so that the relevant comparative statistics do not indicate real-time forecasting accuracy. Then, these simple macroeconomic specifications are examined further by generating forecasts for all the variables in each model for a more truly dynamic forecast. The results of this more realistic exercise show that forecasting accuracy of the representative inflation forecasting models deteriorates (in some ways substantially) and may perform little better than the simplest time series forecasting model.

The simple economic forecasting models examined in this article do not forecast inflation effectively. The single-equation models, based on the intuition of partial equilibrium models that are inherently incomplete, account for observed relationships in the data rather than the underlying forces that drive inflation. The failings of these forecasting models suggest that these intuitions, in their simple form, do not contribute greatly to understanding of the inflation process nor are they effective for forecasting.

The forecasting exercises and the evaluation statistics imply that macroeconomic inflation forecasting models should be evaluated in circumstances that more closely approximate the real-time forecasting problem. The key feature to emphasize, though, about forecasting inflation is that forecasting equations for all variables that are in the inflation forecasting model should be specified. Exercises with static and partially dynamic forecasts help isolate weaknesses in the inflation forecasting models but should not be regarded as indicative of expected forecasting performance because in real-time forecasting there is not the benefit of knowing future values of the explanatory variables as assumed in these two exercises.

The problem of accurately forecasting inflation remains troublesome, especially in light of the concern among policymakers and in financial markets about future inflation. Dynamic systems of equations appear to be a more realistic modeling framework for these inquiries. Fruitful research may focus on estimating dynamic models that are by design more structural and may help uncover the sources of inflation. In terms of optimal inflation forecasting models, research geared toward dynamic specifications of either a structural or nonstructural (but restricted) nature will be most useful for real-time applications.

Notes

1. In fact, the Social Security example ignores the growth rate of wages, which determines the growth rate of Social Security revenues. The example abstracts from other variables relevant for planning Social Security payments like the number of new retirees, emergency survivorship benefits, mortality rate of recipients, and so forth. The net impact of

the inflation forecast error could be higher or lower but would not affect the point of the example.

2. Turning point analysis tends to concentrate on real (inflation-adjusted) output and business cycles, like predicting the next recession, but some research has aimed similar analysis toward inflation. See Roth (1991) and Klein (1986).
3. Modeling using variables that allegedly cause inflation enables the forecaster to interpret the forecast results as being due to movements in a particular variable associated with a set of economic actors (for example, firms, households, and so forth) included in the model. The interpretability of that model may be valued by the forecast consumer, as opposed to forecasts that seem to result from a statistical "black box."
4. The measure squares the forecast errors, thus increasing the weight of large errors, then sums the squared errors, and then takes the square root of the sum.
5. Models truly designed for policymaking need more complete structure than the simple forecasting models presented below. Policymakers would like a fully specified structural model of the economy to generate forecasts of inflation and other variables. This model would attempt to mimic the workings of the actual economy in a small-scale, but internally consistent, model that incorporates government policy actions into the inflation process. The forecasting models below incorporate only a portion of a theory in the design of the model and thus offer an incomplete view of the effects of policy on inflation (see Leeper 1993).
6. McNees (1994) provides analysis of the accuracy of official public sector forecasts (Congressional Budget Office and the Council of Economic Advisors) as well as private business forecasts (Blue Chip Consensus Forecasts). He finds the accuracy of both public and private forecasts of CPI-measured inflation over the past decade to be virtually identical.
7. This statement refers to univariate statistical models.
8. Forecast accuracy criteria may still suggest use of a statistical model; forecast interpretability may be at the cost of less accuracy. An adequate forecasting model need not be unbiased (that is, have on average a zero mean forecast error). The chosen model may, for example, have a mean error that implies overpredicting inflation on average to reduce the likelihood of realizing higher than anticipated inflation. Thus, the measure of forecast accuracy can be designed to fit the criterion most useful to the forecast user.
9. Several existing studies examine more extensively the typical macroeconomic models of inflation forecasting. See, for example, Stockton and Struckmeyer (1989), Stockton and Glassman (1987), Mehra (1988), and Reichenstein and Elliot (1987).
10. Three recent articles in the inflation forecasting literature (Mehra 1988, Stockton and Glassman 1987, and Reichenstein and Elliot 1987) provide useful comparisons of the forecasting effectiveness of a selection of macroeconomic inflation models for a forecast horizon of two years or more. The text describes a selective subset of the models presented in these papers.
11. The estimations below avoid ad hoc empirical additions that may improve the in-sample fit and forecasting statistics for a model but that are not directly related to that model. The results attempt to distinguish between the contribution of "theoretical intuitions" versus other empirical techniques to forecast inflation rather than to determine a "best" model.
12. Frequently, this approach is associated with a Keynesian view of how inflation is generated. Complete derivations of the typical Phillips curve models appear in Glassman and Stockton (1983) and more briefly in Stockton and Glassman (1987) and in Mehra (1988).
13. The following econometric specification is a standard way to formalize the intuition implicit in the "expectations" augmented Phillips curve relationship (accounting for changes in the expected rate of inflation by agents' observations of past inflation rates as seen in the last term):

$$\Delta p_t = c_0 + c_1(y_{t-1} - y_{pt-1}) + c_2(\Delta Y_t - \Delta y_{pt}) + \Sigma_{j=1}^{n} c_{2+j}\Delta p_{t-j} + v_t,$$

where Y_t is the natural log of nominal GDP, y_{pt} is the natural log of potential real GDP, y_t is the natural log of real (observed) GDP, p_t is the natural log of the price level, v_t is the error term, Δ represents the difference operator, and n is the number of lagged values of Δp_{t-j}.

The formulation is a slight variation on the simple GDP gap model. Algebraic manipulation allows the substitution of nominal GDP into the GDP gap formulation. See Mehra (1988). Mehra uses a lagged value of the real GDP gap and the contemporaneous value of the difference between the nominal GDP growth rate and the rate of growth of potential GDP; Stockton and Glassman (1987) use a lagged value of the latter and the contemporaneous value of the former.

14. The limited number of lags of inflation in the estimate is chosen to maximize the degrees of freedom and allow forecast comparisons over the period including 1973 so that both oil shocks are in the out-of-sample forecast period. Results of models with more inflation lags could be noticeably different.
15. As suggested above, it would be reasonable to add other variables to the specification, but the basic premise of this article is to examine the contribution to forecasting accuracy of the basic intuition of the respective models.
16. The results in this article employ only the monetary base as the monetary aggregate; another monetary aggregate would likely produce different results. Similarly, the Phillips curve model employs a potential GDP measure produced by Data Resources, Inc.; another measure of potential GDP would likely produce different results.
17. The inflation forecasting model that results from manipulation of this functional relationship is

$$\Delta p_t = g_0 + \Sigma_{j=1}^{m} g_{1j}\Delta M_{t-j} + \Sigma_{j=1}^{m} g_{2i}\Delta y_{t-j} + \Sigma_{j=1}^{m} g_{3i}\Delta i_{t-j} + u_t,$$

where p_t is the natural log of the price level, M_t is the natural log of the nominal money stock, y_t is the natural log of aggregate output, i_t is the nominal interest rate, Δ is the difference operator, and m is the number of lagged values of the variable.

18. An ARIMA(1, 1,2) model for inflation forecasting is also estimated; the results for that model were sufficiently similar to the simple time series model that they have been excluded in the interest of simplicity.

19. Estimates of the equations over the full sample are available from the author.
20. The values of both the Box-Ljung Q statistic (247.8) and the Durbin-Watson statistic (0.74) indicate autocorrelated errors in the estimated equation. The values of these statistics indicate that the model is poorly specified.
21. The forecast starts from the fourth quarter of 1972 in order to include data observed in the two oil shocks within the forecasting sample period. Models are also estimated over a longer estimation sample (fourth quarter of 1960 to fourth quarter of 1983), and thus a shortened forecast sample, starting in the first quarter of 1984.
22. This comparison, however, is only an approximation because the data used in this study are revised, and real-time forecasting involves forecasting the unrevised values of the data.
23. The article evaluates the forecasting performance of the selected models but is more an exposition on the methods of forecast evaluation than a rigorous search for the optimal inflation forecasting model. Financial models of the term structure have recently been used for inflation forecasting and have been applied to actual data with some success. See Fama (1990) for a clear example. The contributions of these models to inflation forecasting are surveyed in Abken (1993).
24. The forecast statistics could differ noticeably if the estimates were performed over different forecast periods or subsamples.
25. The inflation forecasting studies cited above typically use either eight- or ten-quarter ahead forecasts, leaving relatively few nonoverlapping observations. The longer forecasting horizon may be warranted, though, in specific applications.
26. Stockton and Glassman (1987) produce forecast evaluations using forecasts of the explanatory variables as well as the actual future values of the explanatory variables. Forecasts using actual data on the explanatory variables had noticeably better forecast evaluation statistics. Such an exercise highlights the influence on the forecast errors from forecasts of the explanatory variables on the accuracy of the inflation forecasts.
27. It has been argued that Federal Reserve policy became less tolerant of inflation after 1979.
28. Further inquiry into this issue will be the subject of future research.
29. Recall that the measure of potential GDP is a number that is an available estimate, already forecast a number of periods into the future at the time of the forecast.
30. The specification of the VAR is motivated by the money demand theory described above. Clearly, it would be possible to specify a VAR inspired by a Phillips curve perspective. As used here, VAR refers to the money demand specification.
31. Webb (1984, 1985) has examined extensively the accuracy of VAR inflation forecasting models. For perhaps a more optimistic appraisal of the accuracy of the technique, see Webb (1994).
32. The model illustration below makes this fact more clear:

$$\Delta p_t = g_{010} + \Sigma^m_{j=1} g_{11j} \Delta M_{t-j} + \Sigma^m_{j=1} g_{12i} \Delta y_{t-j} + \Sigma^m_{j=1} g_{13i} \Delta i_{t-j} + \Sigma^m_{j=1} g_{14i} \Delta p_t + u_{1t},$$

$$\Delta M_t = g_{020} + \Sigma^m_{j=1} g_{21j} \Delta M_{t-j} + \Sigma^m_{j=1} g_{22i} \Delta y_{t-j} + \Sigma^m_{j=1} g_{23i} \Delta i_{t-j} + \Sigma^m_{j=1} g_{24i} \Delta p_t + u_{2t},$$

$$\Delta y_t = g_{030} + \Sigma^m_{j=1} g_{31j} \Delta M_{t-j} + \Sigma^m_{j=1} g_{32i} \Delta y_{t-j} + \Sigma^m_{j=1} g_{33i} \Delta i_{t-j} + \Sigma^m_{j=1} g_{34i} \Delta p_t + u_{3t},$$

$$\Delta i_t = g_{040} + \Sigma^m_{j=1} g_{41j} \Delta M_{t-j} + \Sigma^m_{j=1} g_{42i} \Delta y_{t-j} + \Sigma^m_{j=1} g_{43i} \Delta i_{t-j} + \Sigma^m_{j=1} g_{44i} \Delta p_t + u_{4t},$$

where p_t is the natural log of the price level, M_t is the natural log of the nominal money stock, y_t is the natural log of aggregate output, i_t is the nominal interest rate, Δ is the difference operator, m is the number of lagged values of the variable, and u_{it} is the error term associated with equation i.
33. In this application, including the relative price of oil as an explanatory variable increases the root mean squared error of the inflation forecast.
34. The regression estimates for this model are available upon request from the author.

$$\Delta p_t = \alpha_0 + \beta_1 (y_{t-1} - y_{pt-1}) + \beta_2 (\Delta Y_{t-1} - \Delta y_{pt-1}) + \Sigma^4_{j=1} \beta_2 \Delta p_{t-j} + v_t,$$

$$\Delta Y_t = g_0 + \Sigma^4_{j=1} g_{1i} \Delta Y_{t-i} + \Sigma^4_{j=1} g_{2i} \Delta p_{t-j} + \xi_t,$$

$$y_t = \gamma_0 + \gamma_1 T + \Sigma^4_{j=1} \gamma_{2i} y_{t-i} + \Sigma^4_{j=1} \gamma_{3i} \Delta p_{t-j} + \epsilon_t,$$

where Y_t is the natural log of nominal GDP, T is a time trend, y_t is the natural log of real (observed) GDP, p_t is the natural log of the price level, v_t, ϵ_t, and ξ_t are the error terms, Δ represents the difference operator, and n is the number of lagged values of Δp_{t-j}.
35. The Theil U statistic for the dynamic Phillips curve forecasts is 0.896 for the full sample and 0.978 for the shorter sample. In contrast, the Theil U statistics for the VAR (money demand) were 0.980 and 1.368.
36. The potential GDP measure in this model is an exogenous variable, but it is also known that the data series is estimated ex post. As a result, the series may be revised, and through the revision process past data observations of the measure may reflect information not really available at the time period of the observations. The historical data values may reflect important economic information that was not really available, thereby giving some potential advantage to the Phillips curve model. This problem is somewhat analogous to that dealt with in research on the Index of Leading Indicators by Diebold and Rudebusch (1991).
37. The forecasting exercise ignores the additional real-time complication of employing unrevised data for economic aggregates like GDP.

References

Abken, Peter A. "Inflation and the Yield Curve." Federal Reserve Bank of Atlanta *Economic Review* 78 (May/June 1993): 13-31.

Diebold, Fran, and Glen Rudebusch. "Turning Point Prediction with the Composite Leading Index: An Ex Ante Analysis." In *Leading Economic Indicators: New Approaches and Forecasting Records*, edited by Kajal Lahiri and Geoffrey H. Moore, 231-56. Cambridge: Cambridge University Press, 1991.

Elliot, J. Walter, and William Reichenstein. "A Monetary Approach to Measuring Long-Run Inflationary Expectations." *Journal of Economics and Business* 19 (May 1987): 327-38.

Fama, Eugene F. "Inflation, Output, and Money." *Journal of Business* 55, no. 3 (1982): 201-31.

_____. "Term Structure Forecasts of Interest Rates, Inflation, and Real Returns." *Journal of Monetary Economics* 25 (January 1990): 59-76.

Glassman, James E., and David J. Stockton. "An Evaluation of Alternative Price Forecasting Models: Theoretical Considerations." Board of Governors of the Federal Reserve System Working Paper No. 32, December 1983.

Klein, Philip A. "Leading Indicators of Inflation in Market Economies." *International Journal of Forecasting* 2 (1986): 403-12.

Leeper, Eric M. "The Policy Tango: Toward a Holistic View of Monetary and Fiscal Effects." Federal Reserve Bank of Atlanta *Economic Review* 78 (July/August 1993): 1-27.

McNees, Stephen K. "Official and Private Forecasts of the U.S. Economy." Unpublished manuscript, June 1994.

Mehra, Yash P. "The Forecast Performance of Alternative Models of Inflation." Federal Reserve Bank of Richmond *Economic Review* 74 (September/October 1988): 10-18.

Reichenstein, William, and J. Walter Elliot. "A Comparison of Models of Long-Term Inflation Expectations." *Journal of Monetary Economics* 19 (May 1987): 405-25.

Roth, Howard L. "Leading Indicators of Inflation." In *Leading Economic Indicators: New Approaches and Forecasting Records*, edited by Kajal Lahiri and Geoffrey H. Moore, 275-301. Cambridge: Cambridge University Press, 1991.

Stockton, David J., and James E. Glassman. "An Evaluation of the Forecast Performance of Alternative Models of Inflation." *Review of Economics and Statistics* 69 (February 1987): 108-17.

Stockton, David J., and Charles S. Struckmeyer. "Tests of the Specification and Predictive Accuracy of Nonnested Models of Inflation." *Review of Economics and Statistics* 71 (May 1989): 275-83.

Webb, Roy H. "Vector Autoregressions as a Tool for Forecast Evaluation." Federal Reserve Bank of Richmond *Economic Review* 70 (January/February 1984): 3-11.

_____. "Toward More Accurate Macroeconomic Forecasts for Vector Autoregressions." Federal Reserve Bank of Richmond *Economic Review* 71 (July/August 1985): 3-11.

_____. "Forecasts of Inflation from VAR Models." Federal Reserve Bank of Richmond Working Paper 94-8, July 1994.

Review Essay

An Economist's Perspective on History: Thoughts on *Institutions, Institutional Change, and Economic Performance*

Andrew C. Krikelas

History matters. It matters not just because we can learn from the past, but because the present and the future are connected to the past by the continuity of society's institutions. Today's and tomorrow's choices are shaped by the past. And the past can only be made intelligible as a story of institutional evolution. Integrating institutions into economic theory and economic history is an essential step in improving that theory and history.

—Douglass C. North

Institutions, Institutional Change, and Economic Performance, *by Douglass C. North (Cambridge: Cambridge University Press, 1990), 152 pages, $37.95 ($11.95, paper).*

The reviewer is an economist in the regional section of the Atlanta Fed's research department.

For more than two decades, Douglass North's research has focused upon institutions and the important role they play in regulating human social behavior. The ideas that have emerged from this line of inquiry have led North to develop an innovative framework for analyzing economic history, presented most comprehensively in *Institutions, Institutional Change, and Economic Performance*. North acknowledges that this paradigm is not yet fully developed; nevertheless, this body of research earned him a share of the 1993 Nobel prize for economics.[1]

Concerning this award, the *Financial Times* noted, "The Nobel committee's decision to break with tradition and award the prize for work in economic history, rather than economics proper, reflects the growing importance economists attach to the role of social institutions in providing a framework for economic growth."[2] Indeed, the importance of adopting an institutional perspective for analyzing economic events has grown in recent years in the wake of the collapse of the socialist economies in Eastern Europe and the Soviet Union.

While many neoclassically trained economists have been asked to counsel the leaders of nations making the transition from centrally planned to free market economies, much of their advice has proved inadequate if not unhelpful. The main reason for this inadequacy is that most neoclassical economic models are based upon two assumptions: that society's institutional infrastructure is essentially free-market oriented and that it is unchanging. These models, therefore, are not suitable for analyzing situations in which the institutional infrastructure diverges much from the free-market standard or those in which institutional change is the order of the day, and they are particularly inadequate for those in which both conditions prevail simultaneously, as is the case currently in Eastern Europe and the countries of the former Soviet Union. Because institutional economists have not yet developed an alternative behavioral model that accounts for the process of institutional change and its related economic consequences, they too have not been able to offer much specific advice except to criticize the calls for shock therapy made by a number of mainstream economists.

Although many economists remain skeptical that an adequate alternative to the neoclassical paradigm will ever be developed, I would argue that North has already made substantial progress toward this goal. The game theoretic model he describes in *Institutions, Institutional Change, and Economic Performance* is both rigorous and intuitive and deserves a much wider reading among social scientists than it has received to date. The purpose of this review, therefore, is threefold: to stress the book's importance, to explain the essential features of North's model, and to demonstrate the model's usefulness as an analytical tool.

North's View of History

From the outset, it is important to recognize that what North actually offers is a coherent and detailed theory of social history. His methodology allows one to confront, and answer directly, two related questions: What is history, and why does it unfold as it does? In chapter one North introduces a team sport analogy, used throughout the book, that illustrates his theory: Institutions, he states, "are the framework within which human interaction takes place. They are perfectly analogous to the rules of the game in a competitive team sport. That is, they consist of formal written rules as well as typically unwritten codes of conduct that underlie and supplement formal rules, such as not deliberately injuring a key player on the opposing team. And as this analogy would imply, the rules and informal codes are sometimes violated and punishment is enacted. Therefore, an essential part of the functioning of institutions is the costliness of ascertaining violations and the severity of punishment" (4).

From North's perspective history can be characterized as the record of an evolving game. While the game is complex in its detail, its fundamental structure is relatively simple, consisting of just three primary elements: institutions, organizations, and individuals (rules, teams, and players). Although most economists would be inclined to describe the game as an individual sport, North assumes it to be a team competition. He is not saying that individuals do not play an important role in the game. What he does imply is that there are many social, political, and economic objectives that can be obtained by individuals only through team effort. In order for each of the players to succeed as individuals, they must be prepared to behave both cooperatively, in order to produce social wealth, and competitively, in order to consume it. Structurally, then, the game that North describes could be viewed as a cooperative social competition.

At stake in this competition, according to North, are the substantial gains from trade that can be produced by an economic system organized around the dual principles of specialized production and market exchange.[3] These returns, in fact, are exactly the same profit opportunities that fuel the competitive desires of the actors in neoclassical economic models. In contrast to most mainstream economists, who generally assume that players always obey the rules, thereby producing none of the social costs associated with cheating and other antisocial behaviors, North assumes that cheating is inevitable and that referees, therefore, must be hired to resolve conflicts. Thus, in addition to the role played by teams engaged in pure economic activity, North's model contains an explicit role for social and political teams as providers of the services required to maintain orderly economic competition.

The services that social and political teams provide within the context of the game can be divided into three distinct categories (see Table 1). The first set of teams, type I social and political teams, are those that act as referees; their main purpose is to detect rule-breakers and to resolve conflict. Such refereeing services are provided both formally by the state through the police force and judicial system and informally by

Table 1
The Four Types of Organizations That Participate in Cooperative Social Competition

Type of Organization	Objective
Economic Teams	To produce and distribute economic goods and services. These teams are the focal point of the competition.
Type I Social and Political Teams	To detect and punish teams and players who violate established rules and to resolve conflicts.
Type II Social and Political Teams	To teach players the rules.
Type III Social and Political Teams	To lobby for changes in the existing rules.

parents within families, clerics within congregations, and managers within business offices. Although these monitoring services are costly to provide and only imperfectly achieve their goals, they nevertheless play an indispensable role in support of the underlying economic competition.

A second and related set of teams, type II social and political teams, have as their objective the socialization of the individual—the process by which he or she is both taught and persuaded to play by the established rules of the game. Unlike most economists, who would assume that players are born with a complete understanding of the rules, North contends that individuals become aware of the rules only through a lifelong educational process.[4] This process begins within the family with the lessons taught by parents, siblings, and other relatives. Over time, however, the socialization of the individual expands to include lessons taught by many other teams and their players, including educators in schools, clerics in congregations, employers in businesses, and even politicians as leaders of political parties. Because the process by which each person obtains his or her education is unique, it is unlikely that two individuals will have identical perceptions of any given social situation. As a result, North allows for social conflicts to arise not only from willful cheating on the part of teams and players but also from differences in the socialized expectations of the parties engaged in a particular exchange.

The third set of social and political teams participating in this competition is quite different from the first two. While type I and II teams serve to enforce and perpetuate the existing game rules, type III teams seek to rewrite these rules. Such institutional changes may be introduced cooperatively, through prescribed political and legislative channels, or uncooperatively, through the use of force in war or revolution. In each case, however, the social and political teams that initiate these reforms do so on behalf of economic teams—teams whose competitive interests would be served by changes in specific rules. Although the motivations for such changes could stem from many sources, two of the more likely would be a substantial shift in relative prices that finally allows a team to hire previously unaffordable lobbying services and the advent of an innovative technology whose efficient use conflicts with existing rules. In other words, in addition to allowing for the possibility of institutional change and social evolution, North's model identifies the origin to be the entrepreneurial behavior of this group of social and political teams and their players.

Although North does not take the team sport analogy quite as far as this essay has, this elaboration fits well within North's view of history as the record of an evolving game. Unlike in most sports, however, the rules of this game are not fixed but are subject to constant change. Each successive rules change forces all teams and players to reevaluate and modify their corporate and personal strategies, with the result that even small institutional changes may ultimately have a profound impact on the outcome of the competition. In order to understand the outcomes of any given cooperative social competition, it is necessary to collect as much of the following information as possible: (1) a complete set of the rules of the game at the beginning of the competition; (2) a list of the relevant changes to those rules over time; (3) a list of the four types of organizations competing over time, as well as a sense of their corporate strategies; and (4) a list of the individual players, the portfolio of team affiliations they maintain, and the personal strategies they adopt under each set of prevailing rules.

Applying the Model

North's model can be used for a variety of analytical purposes, one of which, of course, is interpreting the events of a specific historical period. In fact, North devotes several chapters of his book to just such an investigation, focusing specifically on the economic history of the Western Hemisphere. In a world in which information and capital flow freely and transactions costs are negligible, neoclassical economic theory predicts that economic convergence will occur, especially in the long run. From this perspective, the divergent economic performances during the last two hundred years of the United States and Canada on the one hand and most of Latin America on the other provide a persistent puzzle.

North solves this puzzle by pointing out that, in a world in which institutions shape the behavior of the teams and players engaged in cooperative social competition, convergence does not necessarily occur. Although two societies might be identically endowed with land, labor, capital, and other productive resources, if their institutional structures are significantly different then they are likely to achieve dramatically different economic results, even in the long run.

For example, the institutions of private property and individual rights that North Americans inherited from their British colonizers established powerful behavioral incentives for private accumulation of both physical and human capital, decentralization of economic and political decision making, and attention to overall productive efficiency. By contrast, the state property rights and central authority that the Spanish conquistadors imposed on colonial Latin America promoted a monopoly over both economic and political resources, encouraging redistributive, rather than productive, activities. North argues that the divergent economic performances observed in the Western Hemisphere during the last two hundred years reflect the very different ways in which these cooperative social competitions have been organized and conducted.

Bridging the Gaps between Academic Disciplines

In addition to being useful for analyzing the history of a particular society, North's game theoretic framework has the potential to serve as a methodological bridge between the diverse branches of the social sciences and humanities. The social sciences—such as economics, political science, sociology, and social psychology—deal with the institutions and functioning of human society and the interrelationship of its members. History, although sometimes categorized as a social science, is generally considered one of the humanities, which can be defined as the branches of learning concerned with human culture and values and which usually include languages, literature, history, mathematics, and philosophy. Although scholars in each of these fields conduct research that would benefit greatly from interdisciplinary discussion, relatively little such communication has occurred because there has been no unifying theory to identify the ways in which these disciplines are related.

North's game theoretic framework goes a long way toward providing just such a unified perspective. It serves as an intuitive lens through which various disciplines can be viewed as one or more aspects of cooperative social competition. History, for example, can be thought of as the record of all outcomes produced by a given social competition, with other humanities focusing specifically on subsets of those outcomes. Within this framework, items of art and literature would be characterized as cultural if they support the socialization efforts of the type I and type II social and political teams and countercultural if they promote the institutional change objectives of type III teams.

North's game theoretic framework has the potential to serve as a methodological bridge between the diverse branches of the social sciences and humanities.

Social scientists, on the other hand, in addition to studying the outcomes of cooperative social competition, also seek to understand the underlying processes that generate these results. While many economists may believe that theirs is the only game in town, North argues quite convincingly that social and political teams play an important and integral role in cooperative social competition. In fact, North's framework clearly suggests that greater interdisciplinary communication among economists, political scientists, sociologists, and social psychologists is required if a better model of cooperative social competition is to be developed.

Conclusion

In order to make sense of history, it is essential to have a model that adequately accounts for the role institutions play in determining social, political, and economic events. North's paradigm is one such model. Although not yet fully specified in terms of the formal mathematics that have become the standard of the mainstream economics literature, North's game theoretic framework nevertheless represents an important step toward a better understanding of why history unfolds as it does.

Because North's behavioral assumptions diverge substantially from those currently adopted by most economists, it is likely to take some time before this new paradigm becomes accepted in the mainstream economics literature. Conversely, because North's methodological approach is essentially that of neoclassical economics, it has been criticized by institutionalists and other noneconomists who are put off by its relative formality.

North himself considers his model an extension, rather than an abandonment, of traditional economic analysis: "Integrating institutional analysis into economics and economic history is redirecting emphasis, but not abandoning the theoretical tools already developed. Redirecting the emphasis entails modifying the notion and implications of rationality, incorporating ideas and ideologies into our analysis, explicitly studying the costs of transacting for the functioning of political and economic markets, and understanding the consequences of path dependence for the historical evolution of economies" (135).

North's new institutional framework offers real promise as a means by which to facilitate interdisciplinary communication among economists, institutionalists, and other social scientists. *Institutions, Institutional Change, and Economic Performance* should become required reading for scholars in all of these fields.

Notes

1. In a 1992 article, North says of this paradigm: "All of this, therefore, does not add up to anything as elegant as a theory. A dynamic theory of economic change is the objective; but what we have so far is a set of definitions and principles and a structure that make up some of the essential scaffolding necessary to a theory of institutional change" ("Institutions and Economic Theory," *American Economist* 36 [Spring 1992]: 6).
2. "Economic Historians Win Nobel," *Financial Times*, October 13, 1993, 7.
3. Some form of market exchange occurs in nearly all economies that engage in specialized production. This is not to say that the market structures in centrally planned and free-market economies are not qualitatively different. However, in both of these economic systems the goods produced by specialized teams are redistributed to other teams and players within society through some form of market exchange.
4. It is at the level of the individual that North's behavioral model diverges most substantially from the neoclassical economic paradigm. While many economists would assume that individuals have the capacity to know and understand completely the rules governing cooperative social competition, North subscribes to Herbert A. Simon's view of bounded rationality, which assumes the individual has only a limited understanding of these formal and informal rules. For a recent discussion of bounded rationality, see chapter 6 of James G. March and Herbert A. Simon, *Organizations* (Cambridge, Mass.: Blackwell, 1993).

FEDERAL
RESERVE
BANK OF
ATLANTA

Public Affairs Department
104 Marietta Street, N.W.
Atlanta, Georgia 30303-2713
(404) 521-8020

Bulk Rate
U.S. Postage
PAID
Atlanta, GA
Permit 292

printed on recycled paper