

Ben S Bernanke: Monetary policy and the stock market - some empirical results

Speech by Mr Ben S Bernanke, Member of the Board of Governors of the US Federal Reserve System, at the Fall 2003 Banking and Finance Lecture, Widener University, Chester, Pennsylvania, 2 October 2003.

The references for the speech can be found on the Board of Governors of the Federal Reserve System's website.

* * *

The ultimate objective of monetary policymakers is to promote the health of the U.S. economy, which we do by pursuing our mandated goals of price stability and maximum sustainable output and employment. However, the effects of our policy instruments, such as the short-term interest rate, on these goal variables are indirect at best. Instead, monetary policy actions have their most direct and immediate effects on the broader financial markets, including the stock market, government and corporate bond markets, mortgage markets, markets for consumer credit, foreign exchange markets, and many others. If all goes as planned, the changes in financial asset prices and returns induced by the actions of monetary policymakers lead to the changes in economic behavior that the policy was trying to achieve. Thus, understanding how monetary policy affects the broader economy necessarily entails understanding both how policy actions affect key financial markets, as well as how changes in asset prices and returns in these markets in turn affect the behavior of households, firms, and other decisionmakers. Studying these links is an ongoing enterprise of monetary economists both within and outside the Federal Reserve System.

The link between monetary policy and the stock market is of particular interest. Stock prices are among the most closely watched asset prices in the economy and are viewed as being highly sensitive to economic conditions. Stock prices have also been known to swing rather widely, leading to concerns about possible "bubbles" or other deviations of stock prices from fundamental values that may have adverse implications for the economy. It is of great interest, then, to understand more precisely how monetary policy and the stock market are related.

In my talk today, I will report the results of research that I have done on this topic with Kenneth Kuttner of the Federal Reserve Bank of New York, as well as the findings of some related work done both within and outside the Federal Reserve System.¹ The views I will express today, however, are my own and not necessarily those of my colleagues on the Federal Open Market Committee (FOMC) or the Board of Governors of the Federal Reserve System.

In our research, Kuttner and I asked two questions. First, by how much do changes in monetary policy affect equity prices? As you will see, we focus on changes in monetary policy that are unanticipated by market participants because anticipated changes in policy should already be discounted by stock market investors and, hence, are unlikely to affect equity prices at the time they are announced. We find an effect of moderate size: Monetary policy matters for the stock market but, on the other hand, it is not one of the major influences on equity prices.

Our second question, both more interesting and more difficult, is, why do changes in monetary policy affect stock prices? We come up with a rather surprising answer, at least one that was surprising to us. We find that unanticipated changes in monetary policy affect stock prices not so much by influencing expected dividends or the risk-free real interest rate, but rather by affecting the perceived riskiness of stocks. A tightening of monetary policy, for example, leads investors to view stocks as riskier investments and thus to demand a higher return to hold stocks. For a given path of expected dividends, a higher expected return can be achieved only by a fall in the current stock price. As we will see, this finding has interesting implications for several issues, including the role of stock prices in transmitting the effects of monetary policy actions to the broader economy and the potential effectiveness of monetary policy in "pricking" putative bubbles in the stock market. I will come back to these issues at the end of my talk. I start, however, with the problem of measuring the effect of monetary policy on the stock market.

¹ Bernanke and Kuttner (2003); <http://home.earthlink.net/~kkuttner/bernanke-kuttner.pdf>

The Effect of Monetary Policy Actions on the Stock Market

Normally, the FOMC, the monetary policymaking arm of the Federal Reserve, announces its interest rate decisions at around 2:15 p.m. following each of its eight regularly scheduled meetings each year. An air of expectation reigns in financial markets in the few minutes before to the announcement. If you happen to have access to a monitor that tracks key market indexes, at 2:15 p.m. on an announcement day you can watch those indexes quiver as if trying to digest the information in the rate decision and the FOMC's accompanying statement of explanation. Then the black line representing each market index moves quickly up or down, and the markets have priced the FOMC action into the aggregate values of U.S. equities, bonds, and other assets.

On occasion, if economic conditions warrant, the FOMC may decide to make a change in monetary policy on a day that falls between regularly scheduled meetings, a so-called intermeeting move. Intermeeting moves, typically agreed upon during a conference call of the Committee, nearly always take financial markets by surprise, at least in their precise timing, and they are often followed by dramatic swings in asset prices.

Even the casual observer can have no doubt, then, that FOMC decisions move asset prices, including equity prices. Estimating the size and duration of these effects, however, is not so straightforward. Because traders in equity markets, as in most other financial markets, are generally highly informed and sophisticated, any policy decision that is largely anticipated will already be factored into stock prices and will elicit little reaction when announced. To measure the effects of monetary policy changes on the stock market, then, we need to have a measure of the portion of a given change in monetary policy that the market had not already anticipated before the FOMC's formal announcement.

Fortunately, the financial markets themselves are a source of useful information about monetary policy expectations. As you may know, the FOMC implements its decisions about monetary policy by changing its target for a particular short-term interest rate, the federal funds rate. The federal funds rate is the rate at which depository institutions borrow and lend reserves to and from each other overnight; although the Federal Reserve does not control the federal funds rate directly, it can do so indirectly by varying the supply of reserves available to be traded in this market. Since October 1988, financial investors have been able to hedge and speculate on future values of the federal funds rate by trading contracts in a futures market, overseen by the Chicago Board of Trade. Investors in this market have a strong financial incentive to try to guess correctly what the federal funds rate will be, on average, at various points in the future. The existence of a market in federal funds futures is a boon not only to investors, such as banks, which want to protect themselves against changes in the cost of reserves, but also to both policymakers and researchers, because it allows any observer to infer from the sale prices of futures contracts the values of the federal funds rate that market participants anticipate at various future dates.² Previous research (Krueger and Kuttner, 1996; Owens and Webb, 2001) has shown that participants in this market collectively do a good job of forecasting future values of the funds rate, efficiently incorporating available information about likely future monetary policy actions.³

By using data from the federal funds futures market, then, it is possible to estimate the value at which financial market participants expect the FOMC to set its target for the federal funds rate on any given date. By comparing this expected value to what the FOMC actually did at each date, we can determine the portion of the Fed's interest rate decision that came as a surprise to financial markets. In our research, Kuttner and I considered all the dates of scheduled FOMC meetings plus all the dates on which the FOMC changed the federal funds rate between meetings, or made intermeeting moves, for the period May 1989 through December 2002, amounting to a total of 131 observations.⁴ For each of these dates, we used the expected value of the federal funds rate as inferred from the futures market to divide the actual change in the federal funds rate on that day into the part that was anticipated by

² The futures contract is based on monthly averages of the federal funds rate, so that some manipulation is needed to obtain the daily expectations of the funds rate used in this paper. See Bernanke and Kuttner (2003) or Kuttner (2001) for further details. Allowing for risk premiums creates another complication; see Sack (2002). I ignore these technicalities here.

³ Other financial instruments, such as eurodollar futures rates, can and have been used to forecast changes in the federal funds rate. Although each of the various alternatives has advantages, Gürkaynak, Sack, and Swanson (2002) find that the federal funds futures rate is the best predictor of monetary policy actions for horizons out to several months.

⁴ The beginning of the sample corresponds to the availability of the futures data. We excluded the observation corresponding to September 17, 2001, the first day of trading following the September 11 terrorist attacks.

the markets and the part that was unanticipated.⁵ So, for example, on November 6, 2002, the Federal Reserve cut the federal funds rate by 50 basis points. (A basis point equals 1/100 of a percentage point, so a 50-basis-point cut equals a cut of 1/2 percentage point.) However, this cut in the federal funds rate was not entirely unexpected; indeed, according to the federal funds futures market, investors were expecting a cut of about 31 basis points, on average, from the Fed at that meeting.⁶ So, of the 50 basis points that the FOMC lowered its target for the federal funds rate last November 6, only 19 basis points were a surprise to financial markets and thus should have been expected to affect asset prices. Note, by the way, that if the Fed had not changed interest rates at all that day, our method would have treated that action as the equivalent of a surprise tightening of policy of 31 basis points because the Fed would have done nothing while the market was expecting an easing of 31 basis points.

To evaluate the effect of monetary policy on the stock market, we looked at how broad measures of stock prices moved on days on which the Fed made unanticipated changes to policy. I can illustrate our method by continuing the example of the Fed's cut in the federal funds rate last November 6. On that day, the broad stock market index we used in our study (the value-weighted index constructed by the Center for Research in Securities Prices at the University of Chicago) rose in value by 0.96 percentage point. Dividing the 96-basis-point gain in the stock market by the 19-basis-point downward surprise in the funds rate, we obtain a value of approximately 5 for the "stock price multiplier" relating policy changes to stock market changes. If this one day were representative, we would conclude that each basis point of surprise monetary easing leads to about a 5-basis-point increase in the value of stocks. Or, choosing magnitudes that might be more helpful to the intuition, we could just as well say that a surprise cut of 25 basis points in the federal funds rate should lead the stock market to rise, on the same day, about 1.25 percentage points--about 120 points on the Dow Jones index at its current value. In fact, applying a formal regression analysis to the full sample from 1989 to 2002, we found a number fairly close to this one, namely, a stock price multiplier for monetary policy of about 4.7. We also found, as expected, that changes in monetary policy that were anticipated by the market had small and statistically unimportant effects on stock prices, presumably because these changes had already been priced into stocks.⁷

Although a stock price multiplier of about five for unanticipated changes in the federal funds rate is certainly not negligible, we should appreciate that unexpected changes in monetary policy account for a tiny portion of the overall variability of the stock market. Unanticipated movements in the federal funds rate of 20 basis points or more are relatively rare (we observed only thirteen examples in our fourteen-year sample). Yet the change of one percent or so in the stock market induced by the typical 20-basis-point "surprise" in the funds rate is swamped by the overall variability of stock prices. For example, over the past five years, the broad stock market has moved one percent or more on about 40 percent of all trading days. Thus, news about monetary policy contributes very little to the day-to-day fluctuations in stock prices.

We explored our empirical results with some care. We noted, for example, that a few of the monetary policy changes in our sample were followed by what seemed to be excessive or otherwise unusual stock market responses. A number of these responses occurred rather recently, during the Fed's series of rate cuts in 2001. The Fed's surprise intermeeting cuts of 50 basis points each on January 3 and April 18 of that year were both greeted euphorically by the stock market, with one-day increases in stock values of 5.3 percent and 4.0 percent, respectively. By contrast, the rate cut of 50 basis points

⁵ That the Federal Reserve has only been formally announcing its policy moves since 1994 added a measure of complexity to our research. Before then, market participants generally did not become aware of the FOMC's policy decisions until those decisions were actually implemented in the market for bank reserves, often the day after the FOMC decision. To the extent possible, we dated the policy change as of the day that the market would have become aware of it, not the day of the decision itself. See the paper for details.

⁶ Investors would not literally expect the Fed to cut the funds rate by 31 basis points, since the Fed usually moves in 25-basis-point increments. An average expectation of a 31-basis-point cut would be consistent with, for example, 62 percent of investors expecting a 50-basis-point and 38 percent expecting no cut.

⁷ In principle, news other than the policy decision might affect the federal funds futures contract during the day, so that the measure of unanticipated policy changes we use here might be a "noisy" one. If so, our approach would underestimate the effect of policy changes on the stock market. However, Poole, Rasche and Thornton (2002, pp. 68-69) perform an analysis that suggests that the mismeasurement may be small in practice. Further confirmation is provided by D'Amico and Farka (2002), who find results similar to ours using ten-minute windows around the announcement; the benefit of a tight window is that the policy announcement is highly likely to dominate movements in the contract over that period.

on March 20, 2001, was received less enthusiastically. Even though the cut was more or less what the futures market had been anticipating, the financial press reported that many equity market participants were "disappointed" that the rate cut hadn't been an even larger 75-basis-point action. In any event, the market lost more than 2 percent that day.

To ensure that our results did not depend on a few unusual observations, or "outliers," we re-ran our regression, omitting the days with the most extreme or unusual market moves. This more conservative analysis led to a smaller estimate of the effect of policy actions on the stock market, a stock price multiplier of about 2.6 rather than 4.7. However, the effect remains quite sharp in statistical terms.⁸

We considered other variations as well. For example, we investigated whether the magnitude of the effect on the stock market of a surprise policy tightening (that is, an increase in interest rates) differs from that of a surprise easing of comparable size. It does not. Yet another experiment consisted of asking whether an unanticipated policy change has a larger effect if it is thought by the market to signal a longer-lasting change in policy. We measured the perceived permanence of policy changes by observing the effects of unanticipated policy changes on the expected federal funds rate three months in the future, as measured by the futures market. The stock market multiplier associated with unanticipated policy moves that are perceived to be more permanent is a bit higher, as would be expected; its value is about 6.⁹

In short, the statistical evidence is strong for a stock price multiplier of monetary policy of something between 3 and 6, the higher values corresponding to policy changes that investors perceive to be relatively more permanent. That is, according to our findings, a surprise easing by the Fed of 25 basis points will typically lead broad stock indexes to rise from between 3/4 percentage point and 1-1/2 percentage points. Incidentally, similar results obtain for stock values of industry groups: We find almost all industry stock portfolios respond significantly to changes in monetary policy, with telecommunications, high-tech, and durables goods industry stocks being the most sensitive to monetary policy news, and energy, utilities, and health care stocks being the least sensitive.¹⁰ These results can be broadly explained by the tendency of each industry group to move with the broad market, or (to use the language of the standard capital asset pricing theory), by their industry "betas."

Why Does Monetary Policy Affect Stock Prices?

It is interesting, though perhaps not terribly surprising, to know that Federal Reserve policy actions affect stock prices. An even more interesting question, though, is, why does this effect occur? Answering this question will give us some insight into how monetary policy affects the economy, as well as the role that the stock market should play in policy decisions.

A share of stock is a claim on the current and future dividends (or other cash flows, such as stock buybacks) to be paid by a company. Suppose, for just a moment, that financial investors do not care about risk. Then only two types of news ought to affect current stock values: news that affects investor forecasts of current or future (after-tax) dividends or news that affects forecasts of current or future short-term interest rates. News that current or future dividends (which I want to think of here as being measured in real, or inflation-adjusted, terms) are likely to be higher than previously expected--say, because the company is expecting to be more profitable--should raise the current stock price. News

⁸ Technically, we removed outlier observations based on their so-called influence statistics, which measure the importance of individual observations to the overall results. Another correction was needed because, in the early part of the sample, particularly between 1989 and 1992, it was not uncommon for intermeeting rate cuts to take place on the same day that the government issued weaker-than-expected reports about employment growth. In such cases, our method cannot distinguish cleanly between the effects of the employment news and the effects of the rate cut itself on the stock market. If we eliminate both the outlier observations and the observations in which employment reports coincided with rate changes, we find the multiplier effect of policy changes on the stock market to be about 3.6 and again statistically significant.

⁹ To focus on policy surprises of longer duration, Rigobon and Sack (2002) derive their measure of the unexpected policy change on the three-month eurodollar deposit rate, rather than the current month's federal funds rate, as in this paper and in Kuttner (2001). Using a methodology that also attempts to correct for two-way causality between the funds rate and asset prices, and data for post-1993 scheduled FOMC meetings and Chairman's testimony dates only, they find comparable though slightly higher values for the effect of monetary policy on the stock market. For example, they find a policy multiplier for the Standard and Poor's 500 index of 7.7. However, when they use data on the federal funds rate futures market to measure policy shocks, Rigobon and Sack find results similar to ours, using their sample and methodology.

¹⁰ Using methods similar to ours, Guo (2002) found that the impact of monetary policy actions on stock prices does not seem to depend on firm size.

that current or future short-term interest rates (also measured in real, or inflation-adjusted, terms) are likely to be higher than previously expected should depress the stock price. There are two essentially equivalent ways of understanding why expectations of higher short-term real interest rates should lower stock prices. First, to value future dividends, an investor must discount them back to the present; as higher interest rates make a given future dividend less valuable in today's dollars, higher interest rates reduce the value of a share of stock. Second, higher real interest rates make investments other than stocks, such as bonds, more attractive, raising the required return on stocks and reducing what investors are willing to pay for them. Under either interpretation, expectations of higher real interest rates are bad news for stocks.

So, to reiterate, in a world in which investors do not care about risk, stock prices should change only with news about current or future dividends or about current or future real interest rates. However, investors do care about risk, of course. Because investors care about risk, and because stocks are viewed as relatively risky investments, investors generally demand a higher average return, relative to other assets perceived to be safer, to hold stocks. Using long historical averages, one finds that, in the United States, a diversified portfolio of stocks has paid 5 to 6 percentage points more per year, on average, than has a portfolio of government bonds. This extra return, known as the risk premium on stocks, or the equity premium, presumably reflects, in part, the extra compensation that investors demand to be willing to hold relatively more risky stocks.¹¹

Like news about dividends and real interest rates, news that affects the risk premium on stocks also affects stock prices. For example, news of an impending recession could raise the risk premium on stocks in two ways. First, the macroeconomic environment is more volatile than usual during a recession, so stocks themselves may become riskier investments. Second, the incomes and wealth of financial investors tend to fall during a downturn, giving them a smaller cushion to support the lifestyles to which they are accustomed (that is, to make house payments and meet other obligations). With less discretionary income and wealth to absorb potential losses, people may become less willing to bear the risks of more volatile financial investments (Campbell and Cochrane, 1999). For both reasons, the extra return that investors demand to hold stocks is likely to rise when bad times loom. With expected dividends and the real interest rate on alternative assets held constant, the expected yield on stocks can rise only through a decline in the current stock price.¹²

We now have a list of three key factors that should affect stock prices. First, news that current or future dividends will be higher should raise stock prices. Second, news that current or future real short-term interest rates will be higher should lower stock prices. And third, news that leads investors to demand a higher risk premium on stocks should lower stock prices.

How does all this relate to the effects of monetary policy on stock prices? According to our analysis, Fed actions should affect stock prices only to the extent that they affect investor expectations about dividends, short-term real interest rates, or the riskiness of stocks. The trick is to determine quantitatively which of these sets of investor expectations is likely to be most affected when the Fed unexpectedly changes the federal funds rate.

To make this determination, we used a methodology first applied by the financial economist John Campbell, of Harvard University, and by Campbell and John Ammer of the Federal Reserve Board staff (Campbell, 1991; Campbell and Ammer, 1993). Putting the details aside, we can describe the basic idea as follows. Imagine that the expectations of stock market investors can be mimicked by a statistical forecasting model that takes relevant current data as inputs and projects estimated future values of aggregate dividends, real interest rates, and equity risk premiums as outputs. In principle, investors could use such a model to make forecasts of these key variables and hence to estimate what they are willing to pay for stocks. Besides a number of standard variables that have been shown to be helpful in making forecasts of such financial variables, suppose we include in the forecasting model our measure of unanticipated changes in the federal funds rate.¹³ That is, we use the

¹¹ The existence of a large equity premium in the past is, of course, no guarantee of an equally large equity premium in the future. The fact that equities are more widely held today than in the past, implying that the risk of equities is more widely shared, is one reason that the equity premium may be lower in the future than it has been in the past.

¹² Of course, a looming recession is likely also to lower expected dividends (bad for stocks) and lower interest rates (good for stocks). Generally, stock prices are a leading indicator, falling ahead of recessions and rising in advance of recoveries (although with many false signals).

¹³ Variables used in our forecasting model, besides the excess return on stocks, the one-month real interest rate, and the unanticipated change in the funds rate, include the relative bill rate (defined as the three-month Treasury bill rate minus its

information contained in these unanticipated changes in making our forecasts of future dividends, interest rates, and risk premiums.

Now we can consider the following thought experiment. Suppose we have run our computer model, made our forecasts, and inferred the appropriate values for stocks. But then we receive news that the Fed has unexpectedly raised the federal funds rate by 25 basis points. Based on our forecasting model, by how much would that information change our previous forecasts of future dividends, interest rates, and risk premiums? The answer to this question clarifies the channel by which monetary policy affects stock prices. If we were to find, for example, that the news of an unexpected increase in the funds rate significantly changed the forecast of future dividends but did not much affect the forecasts of interest rates or risk premiums, then we could conclude that monetary policy affects stock prices primarily by affecting investor expectations of future dividends. By contrast, if news of the policy action changed the model forecasts for real interest rates but did not change our forecasts for the other two variables, we would decide that unanticipated policy actions affect stock prices primarily by influencing the interest rates expected by stock investors.

What we actually found when conducting this statistical experiment was quite interesting. It appears that, for example, an unanticipated tightening of monetary policy leads to only a modest change in forecasts of future dividends and to still less of a change in forecasts of future real interest rates (beyond a few quarters). Quantitatively, according to our methodology, the most important effect of a policy tightening is on the forecasted risk premium. Specifically, an unanticipated tightening of monetary policy raises expected risk premiums on stocks for a protracted period. For a given expected stream of dividend payouts and real interest rates, the risk premium and hence the return to holding stocks can only rise if the current stock price falls.

In short, our analysis suggests that an unanticipated monetary tightening lowers stock prices only to a small extent by lowering investor expectations about future dividend payouts, and by still less by raising expected real interest rates. The most powerful effect of an unanticipated monetary tightening is to increase the perceived risk premium on stocks, either by increasing the riskiness of stocks, by reducing people's willingness to bear risk, or both. Reduced willingness of investors to hold relatively more risky stocks drives down stock prices.

Our analysis does not explain precisely how monetary policy affects risk, but we can make reasonable conjectures. For example, tighter monetary policy may raise the riskiness of shares themselves by raising the interest costs and weakening the balance sheets of publicly owned firms (Bernanke and Gertler, 1995). In the macroeconomy more generally, by reducing spending and economic activity, tighter money raises the risks of unemployment or bankruptcy faced by individual households or firms. In each case, tighter monetary policy increases risk by reducing financial buffers or otherwise increasing the vulnerability of individuals or firms to future shocks to the economy.

Implications of the Results for Monetary Policy

So far I have discussed two principal conclusions from the empirical analysis: First, the stock price multiplier of monetary policy is between 3 and 6—in other words, an unexpected change in the federal funds rate of 25 basis points leads, on average, to a movement of stock prices in the opposite direction of between 3/4 percentage point and 1-1/2 percentage points. Second, the main reason that unanticipated changes in monetary policy affect stock prices is that they affect the risk premium on stocks. In particular, a surprise tightening of policy raises the risk premium, lowering current stock prices, and a surprise easing lowers the risk premium, raising current stock prices.

What implications do these results have for our broader understanding and for the practice of monetary policy? I will briefly discuss two issues: first, the role of the stock market in the transmission of monetary policy changes to the economy; and second, the efficacy of monetary policy as a tool for controlling stock market "bubbles."

A long-held element of the conventional wisdom is that the stock market is an important part of the transmission mechanism for monetary policy. The logic goes as follows: Easier monetary policy, for example, raises stock prices. Higher stock prices increase the wealth of households, prompting

12-month moving average), the change in the bill rate, the smoothed dividend-price ratio, and the spread between 10-year and one-month Treasury yields.

consumers to spend more--a result known as the wealth effect. Moreover, high stock prices effectively reduce the cost of capital for firms, stimulating increased capital investment. Increases in both types of spending--consumer spending and business spending--tend to stimulate the economy.

This simple story can be elaborated somewhat in light of our results. It is true, as I have discussed, that an easier monetary policy raises stock prices, whereas a tighter policy lowers them. However, easier monetary policy not only raises stock prices; as we have seen, it also lowers risk premiums, presumably reflecting both a reduction in economic and financial volatility and an increase in the capacity of financial investors to bear risk. Thus, our results suggest that easier monetary policy not only allows consumers to enjoy a capital gain in their stock portfolios today, but it also reduces the effective amount of economic and financial risk they must face. This reduction in risk may cause consumers to trim their precautionary saving, that is, to reduce the amount of income that they put aside to protect themselves against unforeseen contingencies. Reduced precautionary saving in turn implies more spending by households. Thus, the reduction in risk associated with an easing of monetary policy and the resulting reduction in precautionary saving may amplify the short-run impact of policy operating through the traditional channel based on increased asset values. Likewise, reduced risk and volatility may provide an extra kick to capital expenditure in the short run, as firms are more likely to undertake investments in new structures or equipment in a more stable macroeconomic environment.¹⁴

A second issue concerns the role of monetary policy in the management of large swings in stock values, or "bubbles." In an earlier speech (Bernanke, 2002), I gave a number of reasons why I believe that using monetary policy--as opposed to microeconomic, prudential policies--is not a good way to address the problem of asset-market bubbles. These included the difficulty of identifying bubbles in advance; the questionable wisdom, in the context of a free-market economy, of setting up the central bank as the arbiter of asset values; the problem that arises when a bubble occurs in only one asset class rather than in all asset classes; and other reasons. A major concern that I have about the bubble-popping strategy, however, is that attempts to bring down stock prices by a significant amount using monetary policy are likely to have highly deleterious and unwanted side effects on the broader economy. The research I have described today allows me to address this issue more concretely. Here I will make just two points.

First, this research suggests that relatively small changes in monetary policy would not do much to curb a major overvaluation in the stock market. As we have seen, a surprise tightening of 25 basis points should be expected to lower stock prices by only a little more than 1 percent, which, as already noted, is a trivial movement relative to the overall variability of the stock market. It would not be appropriate to extrapolate these results to try to estimate how much tightening would be needed to correct a substantial putative overvaluation in stock prices, but it seems clear that a light tapping of the brakes will not be sufficient. What we can say is that the necessary policy move would have to be quite large--many percentage points on the federal funds rate--and we would be highly uncertain about its magnitude or its ultimate effects on stock prices and the economy.^{15 16}

¹⁴ There is a bit more to this analysis. An additional complexity arises from the fact that, although easier monetary policy allows consumers to enjoy a capital gain in their stock portfolios today, it also "takes back" some of that gain, so to speak, by affording shareholders a lower rate of return on their holdings, on average, in subsequent periods. Research by Sydney Ludvigson and Martin Lettau of New York University and Charles Steindel of the Federal Reserve Bank of New York (Ludvigson, Steindel, and Lettau, 2002; Lettau and Ludvigson, 2001) suggests that, because the gain in share prices induced by a monetary easing is partly transitory, consumers will not increase their spending in response to stock price changes induced by monetary policy as much as they will in response to stock price changes induced by other factors. The estimates in our paper suggest that this differential effect will be relatively small, however. Also, to the extent that the capital gains induced by monetary policy are perceived as partly transitory, the short-run response of investment spending will be strengthened, as firms prefer to invest while stock prices remain high; see Lettau and Ludvigson, 2002, for evidence. In short, if changes in stock values induced by monetary policy are perceived as relatively more transitory, the effects of policy will be concentrated more on investment spending and less on consumption spending than the conventional wisdom suggests.

¹⁵ Greenspan (2002) notes several episodes in which increases in the federal funds rate of several hundred basis points did not materially slow stock appreciation. He argues that "such data suggest that nothing short of a sharp increase in short-term rates that engenders a significant economic retrenchment is sufficient to check a nascent bubble." The late Fischer Black once defined an efficient stock market as one in which prices are between half and double fundamental values; if Black's view is to be believed, then identifiable deviations of prices from fundamentals would have to be quite large indeed.

¹⁶ Implicitly I am considering here the case of a central bank that responds only sporadically to stock prices, in those situations in which it perceives a bubble to be forming. Irregular deviations from a policy rule focused on output and inflation seem appropriately modeled as unanticipated movements in policy. An alternative policy strategy would be to incorporate regular

Second, we have seen that monetary tightening reduces stock prices primarily by increasing the risk premium for holding stocks, as opposed to raising the real interest rate or lowering expected dividends. The risk premium for stocks will rise only to the extent that broad macroeconomic risk rises, or that people experience declines in income and wealth that reduce their ability or willingness to absorb risk (Campbell and Cochrane, 1999). This evidence supports the proposition that monetary policy can lower stock values only to the extent that it weakens the broader economy, and in particular that it makes households considerably worse off. Indeed, according to our analysis, policy would have to weaken the general economy quite significantly to obtain a large decline in stock prices.

Conclusion

I have reported today on empirical work, by my coauthor and me as well as by others, about the links between monetary policy and the stock market. I have only touched on a large literature, and I apologize to the many researchers whose work I have not been able to describe today. But I hope that I have given you a flavor of how empirical research can help us to refine our understanding of how monetary policy works and how policy should be conducted.

reactions to stock values into the systematic part of the monetary policy reaction function. That strategy has some advantages, but it has the important disadvantage that it does not discriminate between fundamental and nonfundamental sources of changes in stock values. Bernanke and Gertler (2001) present simulations showing that such a strategy is unlikely to be beneficial in terms of overall macroeconomic stability.